

HoloHisto: End-to-end Gigapixel WSI Segmentation with 4K Resolution Sequential Tokenization

Yucheng Tang, NVIDIA Corporation, Santa Clara, CA

Yufan He, NVIDIA Corporation, Santa Clara, CA

Vishwesh Nath, NVIDIA Corporation, Santa Clara, CA

Pengfeig Guo, NVIDIA Corporation, Santa Clara, CA

Ruining Deng, Department of Computer Science, Vanderbilt University, Nashville, TN

Tianyuan Yao, Department of Computer Science, Vanderbilt University, Nashville, TN

Quan Liu, Department of Computer Science, Vanderbilt University, Nashville, TN

Can Cui, Department of Computer Science, Vanderbilt University, Nashville, TN

Yuechen Yang, Department of Computer Science, Vanderbilt University, Nashville, TN

Mengmeng Yin, Department of Pathology, Microbiology and Immunology, Vanderbilt University, Nashville, TN

Ziyue Xu, NVIDIA Corporation, Santa Clara, CA

Holger Roth, NVIDIA Corporation, Santa Clara, CA

Daguang Xu, NVIDIA Corporation, Santa Clara, CA

Haichun Yang, Department of Pathology, Microbiology and Immunology, Vanderbilt University, Nashville, TN

Yuankai Huo, Department of Computer Science, Vanderbilt University, Nashville, TN

Abstract

In digital pathology, the traditional method for deep learning-based image segmentation typically involves a two-stage process: initially segmenting high-resolution whole slide images (WSI) into smaller patches (e.g., 256×256 , 512×512 , 1024×1024) and subsequently reconstructing them to their original scale. This method often struggles to capture the complex details and vast scope of WSIs. In this paper, we propose the holistic histopathology (HoloHisto) segmentation method to achieve end-to-end segmentation on gigapixel WSIs, whose maximum resolution is above $80,000 \times 70,000$ pixels. HoloHisto fundamentally shifts the paradigm of WSI segmentation to an end-to-end learning fashion with 1) a large (4K) resolution base patch for elevated visual information inclusion and efficient processing, and 2) a novel sequential tokenization mechanism to properly model the contextual relationships and efficiently model the rich information from the 4K input. To our best knowledge, HoloHisto presents the first holistic approach for gigapixel resolution WSI segmentation, supporting direct I/O of complete WSI and their corresponding gigapixel masks. Under the HoloHisto platform, we unveil a random 4K sampler that transcends ultra-high resolution, delivering 31 and 10 times more pixels than standard 2D and 3D patches, respectively, for advancing computational capabilities. To facilitate efficient 4K resolution dense prediction, we leverage sequential tokenization, utilizing a pre-trained image tokenizer to group image features into a discrete token grid. To assess the performance, our team curated a new kidney pathology image segmentation (KPIs) dataset with WSI-level glomeruli segmentation from whole mouse kidneys. From the results, HoloHisto-4K delivers remarkable performance gains over previous state-of-the-art models.

Introduction

Digital pathology, a rapidly evolving field of medical vision research, has seen a transformative advancement with large vi-

sion models (LVMs) [1, 26, 33]. This significant advent created new demands for high-quality perception, which is crucial for microscopic (e.g., whole slide) image computing. However, current models are limited to the capability of dissecting and interpreting small pre-defined patches within images [8, 10]. Typically, pre-processed tiles are confined to dimensions of 512×512 pixels or resampled to smaller dimensions of 224×224 defined by some predominating frameworks [5, 12, 28], which restrict the scopes of tissue details that can be captured. The absence of rich information hinders the model's performance, particularly impacting tasks of detecting small objects and dense prediction [10, 14]. For instance, the detection and segmentation of complete medullas under kidney WSI will degrade by more than 10% in DSC by using a height and width of 512 or completely fail without patching pre-defined ROI [11]. This scalability, especially when dealing with gigapixel whole slide image (WSI), remains a bottleneck in comprehensive and efficient computing analysis. To date, there are no established gold standard datasets for segmenting gigapixel WSIs, resulting in a lack of comprehensive end-to-end methods in this histopathology research.

To include more information, a higher resolution is necessary as shown in Fig. 1. Nevertheless, modeling ultra-high definition (UHD) images (e.g. beyond 4K resolution) is extremely challenging [9, 13, 30]. High-resolution dense prediction requires a balance of strong context information extraction and model efficiency [4]. The computational cost of convolutional and transformer models, despite their significant benefits, has quadratic increases in the demand for computational resources [31]. This presents a critical scaling challenge for processing whole slide images. Therefore, high-quality image dense prediction requires models capable of understanding both global composition and locality of interactions at a compression rate.

In this work, we propose the HoloHisto framework, debuting the holistic approach to redesign histopathology image segmentation with three key features:




Figure 1: The field of view provided by different resolutions, such as the conventional 512×512 patch (indicated by a red box) and 256×256 , reveals only a limited range of details within the target tissue structures. In contrast, ultra-high resolution (4K) images offer a more comprehensive view of interrelations and can serve as a foundational visual constitute for WSI analysis.

The Holistic Approach: We developed an end-to-end workflow for training and inferencing gigascale WSI, introducing a novel learning paradigm to the field of WSI analysis. Holo-Histo is designed to handle inputs and outputs of any size, regardless of whether they are (WSIs) or smaller patches. By leveraging cuCIM, our dataloader facilitates real-time reading of WSIs at various magnification levels and supports random foreground patching, tiling, or augmentation, enhancing the flexibility and efficiency of our approach. Our approach is capable of dynamically creating datasets online from one or multiple WSIs, potentially comprising an unlimited number of images during training. This approach does not depend on pre-defined cropping strategies, offering a more flexible and scalable solution for training models on large-scale datasets. In the inference stage, it can generate the corresponding gigapixel output.

Architecture: We design an efficient backbone tailored for segmenting UHD images. First, we employ the sequential tokenizer for learning discrete visual tokens from perceptually rich constituents, streamlining towards 4K resolution dense prediction. Second, to model the long discrete tokens from these UHD images, we propose to use a two-stage ViT architecture that incorporates multi-scale attention [4], which uses ReLU linear attention instead of the inefficient Softmax attention.

Data: As a significant effort to improve gigapixel WSI computing, our pathologists addressed the critical gap in the availability of imaging data. We present Kidney Pathology Image Segmentation (KPIS), the dataset that facilitates the diagnosis and treatment of chronic kidney disease (CKD). Annotations are performed at the WSI level, serving as a foundation benchmark for developing cutting-edge image segmentation technologies.

In summary, this paper explores a new learning paradigm of WSI segmentation: 1) HoloHisto framework that is capable of paralleling tile processing with direct WSI I/O, 2) scalable segmentation backbone with sequential tokenizer for ultra-high resolution images, and 3) gigapixel WSI annotation dataset as a foundational benchmark.

Related Works

Pathology Segmentation. Recent advances in deep learning with CNN and transformers [15, 18] achieve significant improvement in the field of pathology segmentation. Several works were proposed to address the challenges of microscopic imaging data, including H&E stained pathology images, fluoresce data, or other cell imaging modalities [3, 22]. Numerous datasets for cell and tissue segmentation, including MoNuSeg [25] and NEPTUNE [2], are available for identifying a variety of glomerular structures. In addition, instance segmentation is developed in the general cell imaging domain [27]. However, most current approaches focus on analyzing local tiles at a uniform magnification level, including nuclei, glomeruli, or tubules. This results in a notable gap in the segmentation of disease-related regions across entire whole slide images. Despite limited exploration or established efficacy in the field, we introduce a segmentation dataset and methodology designed for comprehensive WSI segmentation.

Foundation Vision Models. Inspired by the achievement of large language models (LLMs) [13, 32], many endeavors [1, 17] have been made to develop foundation vision models. With the development of transformer or state space models [16], sequence modeling became the de facto way for modeling visual sentences [1], which enabled the uniform modeling of various vision tasks. In this work, we explore large vision models (LVM) for digital pathology with two key features: (1) a pre-trained vector quantized generative adversarial networks (VQGAN) [13] that enables scalable tokenization for the ultra-high-resolution image at a compression rate; (2) an efficient multi-scale attention module for long sequence representation learning.

Approach

In this work, we propose a holistic framework for segmenting gigapixel WSI. In addition, to model ultra-high resolution representation for dense prediction, we propose a model architecture for high-quality perception learning: 1) use the sequence tokenization for learning 4K visual parts at compression scale; 2) train ViT blocks with linear multi-scale attention. We summarize our approach in Fig 2.

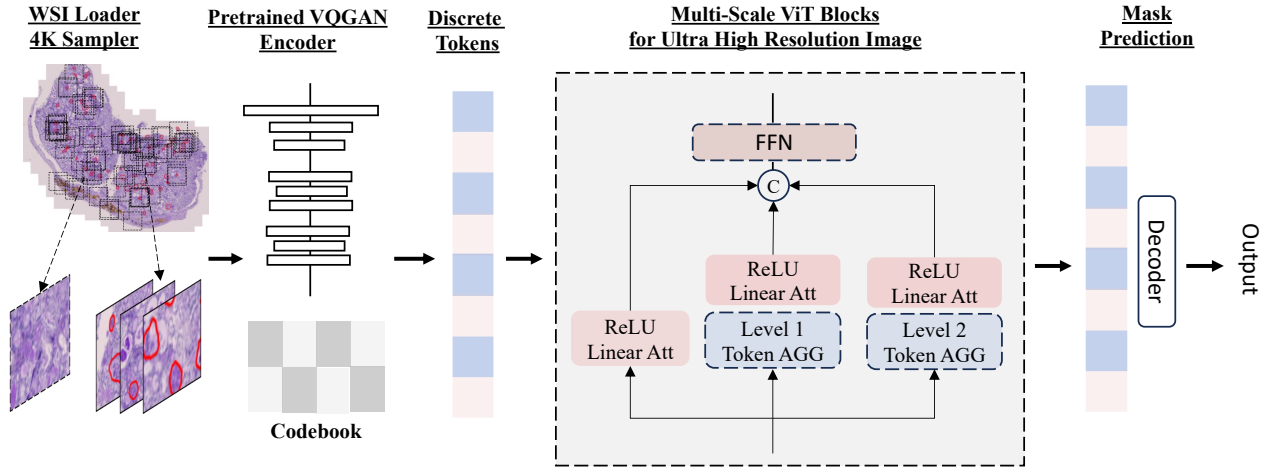


Figure 2: **HoloHisto-4K backbone.** To enable scalable encoding of ultra-high-resolution images (4K), our approach uses a pre-trained convolutional VQGAN to learn the from visual parts into discrete tokens, this design follows autoregressive LVM that enables compression while retaining high perception quality. We employed multi-scale linear attention as an efficient way to capture long discrete visual tokens for high-resolution dense prediction.

Sequence Tokenization

To enable scalable modeling of ultra-high-resolution images while circumventing the quadratic increase in complexity associated with the scan-line order of patches, a discrete approach is essential. This method should efficiently encode visual elements and enable the sampling of high-quality perceptual representations from a probabilistic distribution. Inspired by neural discrete representation learning [32] and Vector Quantised (VQGAN) [13], we employ an image tokenizer. This tokenizer maps input images to a semantic, discrete token codebook through a quantization layer. This technique captures the rich constituents of images, effectively reducing the expression length of original 4K resolution images. Consequently, it enables efficient modeling of global interrelations.

Let the given UHD input be denoted by x , which exists in the space $\mathbb{R}^{H \times W \times 3}$. This image can be decomposed into a grid of codebook vectors z_{enc} , within the domain $\mathbb{R}^{h' \times w' \times d_z}$. Here, d_z represents the number of dimensions for each code. We approximate a given image x by $\hat{x} = G(z_q)$.

To obtain z_q , we start with the encoding $\hat{z} = E(x)$, which resides in the space $\mathbb{R}^{h' \times w' \times d_z}$. Following this, we apply an element-wise quantization $q(\cdot)$ to each spatial code \hat{z}_{ij} within \mathbb{R}^{n_z} , aligning it with its nearest entry z_k . The process is formulated as:

$$z_q = q(\hat{z}) := \arg \min_{z_k \in Z} \|\hat{z}_{ij} - z_k\| \quad (1)$$

where $z_q \in \mathbb{R}^{h' \times w' \times d_z}$

Linear Multi-Scale Attention

High-resolution dense prediction models require strong representation learning capability with good efficiency. Instead of widely used Softmax attention [12], ReLU attention [23] provides linear complexity, which offers the flexibility of multi-scale modules for high-resolution dense prediction. Following the efficient ViT [4] design, we make transformer blocks consisting of 2-stage multi-scale ReLU attention and FNN layers. The 3 hier-

archical multi-scale ReLU attention can be expressed as:

$$A_i = \frac{\text{ReLU}(Q_i) \left(\sum_{j=1}^N \text{ReLU}(K_i)^T V_j \right)}{\text{ReLU}(Q_i) \left(\sum_{j=1}^N \text{ReLU}(K_j)^T \right)} \quad (2)$$

The calculations for the terms $\left(\sum_{j=1}^N \max(0, K_j)^T V_j \right)$ and $\left(\sum_{j=1}^N \max(0, K_j)^T \right)$ need to be performed only once.

HoloHisto: End-to-end framework

The complete pipeline of training and inference is demonstrated in Fig 3.

Training Paradigm. In prior studies [7, 10], pathology image training has been performed using pre-cropped patches of a fixed size over selected regions of interest (ROIs). This offline pre-processing, used before the training and inference phases for gigapixel images, results in the model repeatedly learning from the same patches in each epoch. In this work, we introduce a random sampling paradigm for digital pathology image loader based on the cuCIM¹ multidimensional processing unit. During the training phase, a foreground mask is created using a thresholding technique. Subsequently, we randomly extract ROIs at a 4K resolution from the identified foreground areas. The dataloader then compiles a dataset from one or several whole slide images (WSIs). As the number of training epochs increases, the framework is capable of sampling a virtually "unlimited" number of patches from the WSIs.

Inference with WSI. During the inference stage, HoloHisto is capable of processing the entire WSI. The dataloader seamlessly reads the designated magnification level and isolates the foreground regions through thresholding. Subsequently, HoloHisto performs the foreground tiling with or without overlap, and loads individual tiles into one-dimensional GPU buffers, then positions

¹<https://developer.nvidia.com/cucim-for-image-io-processing>

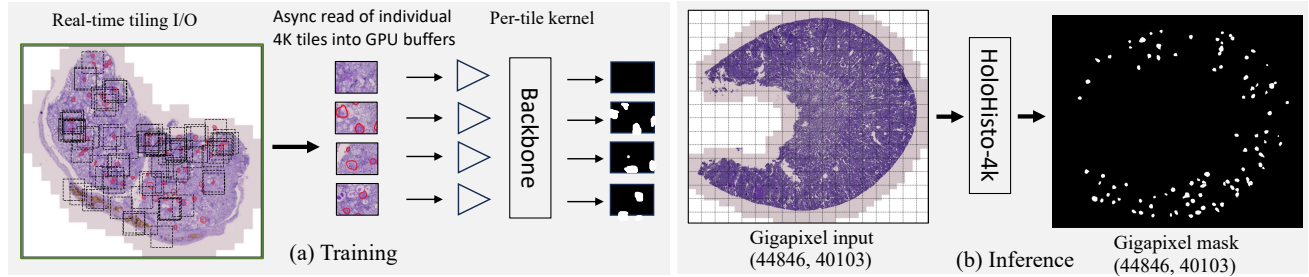


Figure 3: The holistic approach for gigapixel WSI segmentation. HoloHisto takes the entire kidney slide image as input, it supports real-time reading of multi-magnification levels and tile sampling. In the inference stage, HoloHisto presents a global gigapixel segmentation mask as output.

Table 1. Segmentation results of the KPIS dataset on testing cases.

Except HoloHisto, evaluations of WSI on baseline models are performed using predictions from non-overlapping tiles, which are mapped back to the WSI space. Dice similarity coefficient scores (%) are reported. The difference between the reference (Ref.) method and benchmarks for WSI is statistically evaluated by Wilcoxon signed-rank test.

Method	56Nx	DN	NEP25	Normal	All		
					4K Patch-wise	WSI	Statistic.
U-Nets [21]	88.29	87.14	84.91	90.12	87.62	62.19	$p < 0.01$
UNETR [19]	88.96	87.47	85.81	88.46	87.68	64.25	$p < 0.01$
SegFormer [34]	88.58	89.86	89.97	90.33	89.69	65.01	$p < 0.01$
DeepLabV3 [6]	88.41	89.84	89.90	89.91	89.52	68.22	$p < 0.01$
SwinUNETR-V2 [20]	89.04	91.02	89.05	90.08	89.80	71.58	$p < 0.01$
SAM-ViT-B [24]	89.98	90.19	90.13	90.23	90.13	75.19	$p < 0.01$
SAM-ViT-H [24]	90.93	91.04	90.18	90.05	90.55	77.24	$p < 0.01$
HoloHisto-4K	93.77	92.45	94.81	94.12	93.79	84.54	Ref.

them correctly within a pre-allocated GPU array until predictions have been made for all tiles. Finally, the predicted masks for each 4K tile can be allocated back onto the WSI space.

Experiments and Results

Datasets

Kidney Pathology Image Segmentation (KPIs). The KPIs challenge cohort includes 60 high-resolution WSIs of whole mouse kidneys derived from 20 rodents, including three CKD disease models in addition to normal kidneys. Tissues were stained with Periodic Acid-Schiff (PAS) and captured at $40\times$ magnification using the Leica SCN400 Slide Scanner. This diverse cohort allows for comprehensive analysis across different CKD models and normal conditions, providing a rich dataset for advancing research in renal pathology image segmentation. These WSIs are derived from four different groups of mouse models, each representing a different condition or stage of CKD. More information about KPIs studies and annotations is in the supplementary material.

NEPTUNE [2]. The public dataset consists of 1751 Region of Interest (ROI) images that are extracted from 459 Whole Slide Images (WSIs) from 125 patients diagnosed with Minimal Change Disease. These images underwent manual segmentation for six structurally normal pathological features. Each image is at 3000×3000 resolution.

Experiments

We conduct a comparative study of the proposed HoloHisto in two datasets: KPIS and the publicly available tissue segmentation NEPTUNE [2]. For the evaluation of KPIS dataset, we present comparisons among conventional tile-based segmentation

Table 2. NEPTUNE segmentation performance. HoloHisto experiments are conducted under 3000×3000 according to the dataset’s original resolution. Dice similarity coefficient scores (%) are reported.

Metrics	DT	PT	CAP	TUFT	VES	PTC	Average
SAM ViT-B Binary [24]	79.82	85.58	93.58	93.09	83.29	74.90	86.09
UNet multi-scale [29]	81.11	89.85	96.70	96.66	85.03	77.19	87.76
DeepLabV3 [6]	81.05	89.90	96.77	96.69	85.35	78.04	87.97
SwinUNETR-V2 [20]	81.10	89.02	96.74	85.29	85.33	78.68	86.03
SAM ViT-B multi-scale [24]	81.38	89.01	96.90	96.79	85.25	78.58	87.99
SAM ViT-H multi-scale [24]	81.40	90.58	97.00	96.95	85.91	79.05	88.48
HoloHisto-4K	82.14	90.88	97.06	96.99	86.11	79.93	88.85

Table 3. Ablation Study of approach modules. Dice score of 4k tiles and WSIs are reported for ultra-high resolution encoder and WSI handling configurations, respectively.

	Ultra-high Resolution Encoder			WSI Handling		
	Tokenizer	ReLU Att		Sampler	Resolution	
Linear Proj+MHSA	90.45	Linear Proj+ReLU Att	91.85	tile sampler	80.87	2K 79.41
Tokenizer+MHSA	92.94	Tokenizer+ReLU Att	93.79	rand sampler	84.54	4K 84.54

frameworks and calculated metrics in each 4K patch. In addition, to show the effectiveness of the WSI-level segmentation, we compute the Dice score on the entire WSI foreground. We re-trained baseline models including U-Nets [29], UNETR [19], swinunetr-V2 [20], SegFormer [34], and SAM variants [24]. We choose these methods based on groups of CNN, transformer and foundational backbones.

Ultra-high Resolution Analysis

KPIs. Table 1 shows the quantitative result for the binary segmentation task of high-resolution images in the KPIS dataset. We compared HoloHisto to various baselines including CNN, and Transformer-based methods. We evaluated the metrics in two formats: 1) calculate Dice scores under the 4K resolution patch, HoloHisto is trained and inferred in 4K patch, baseline methods are performed in 1024×1024 , which is the best scale for SAM and others. HoloHisto consistently outperforms state-of-the-art pathology segmentation backbone. Along with the ablation study on resolution, we observe the higher resolution patch dimensions, the larger margin is obtained from HoloHisto, indicating the effectiveness of the high-quality perception modeling brought by the tokenizer and efficientViT. In Table 3, we show the ablative experiments result of components design of sequence tokenizer and ReLU linear attention compared to linear projection and multi-head self-attention (MHSA) in vanilla ViT.

NEPTUNE. We conducted additional experiments on the existing public dataset NETUNE. Among 6 different scales of tissues, HoloHisto surpasses baseline models consistently. HoloHisto experiments are performed in 3000×3000 at its largest resolution from the data source, baselines used 1024×1024 sliding window




Figure 4: Kidney pathology segmentation qualitative results. We show the comparisons of three different approaches from CNN, ViT-based SAM and our HoloHisto. The right column shows HoloHisto’s capability of outputting the entire WSI segmentation.

inference from their best training strategy. The Dice scores are reported in Table 2.

End-to-end Prediction

Comparison with pre-tiling In Table 3, the WSI handling section shows the results of using a 4K random sampler with cuCIM and MONAI dataloader versus the tiling strategy, we observe a larger margin of improvement using the end-to-end framework. The visualization of complete WSI is shown in Fig. 4 right panel.

Discussion and Conclusion

This work tackles the fundamental task of segmenting histopathology images, a task that formerly relied on complex pipelines and was restricted to the analysis of small patches. We propose a holistic approach to segment gigapixel images with direct WSI I/O. To model the ultra-high resolution images within loaded WSI, we propose to use a sequential tokenizer, which encodes images as a composition of perception parts and thereby avoids the quadratically increased complexity. In addition, we evaluate the linear ReLU multi-scale attention instead of the Soft-max attention for 4K UHD image tokens. In experiments, we exhibit the first WSI-level segmentation via a 4K image patch sampler and show the effectiveness and capability of HoloHisto-4K by outperforming state-of-the-art approaches. Towards the development of cutting-edge computational research, we also provide the gold-standard pathologist annotated dataset as a WSI segmentation benchmark.

Limitation. It is important to note that we employed the natural image pre-trained sequence tokenizer, where the learned codebook is not for histopathology images. It is still rather challenging to achieve pathology LVM, limiting the model performance and application to WSI analysis. Therefore, we will continue to explore generalist models for pathology vision tasks.

Acknowledgment

This research was supported by NIH R01DK135597 (Huo), DoD HT9425-23-1-0003 (HCY), NSF 2434229 (Huo) and KPMP Glue Grant. This work was also supported by Vanderbilt

Seed Success Grant and Vanderbilt-Liverpool Seed Grant. This research was also supported by NIH grants R01EB033385, R01DK132338. We extend gratitude to NVIDIA for their support by means of the NVIDIA hardware grant.

References

- [1] Bai, Y., Geng, X., Mangalam, K., Bar, A., Yuille, A., Darrell, T., Malik, J., Efros, A.A., Sequential modeling enables scalable learning for large vision models, arXiv preprint arXiv:2312.00785 (2023).
- [2] 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., et al., Digital pathology evaluation in the multicenter nephrotic syndrome study network (neptune), Clinical journal of the American Society of Nephrology: CJASN 8(8), 1449 (2013).
- [3] Bueno, G., Fernandez-Carrobles, M.M., Gonzalez-Lopez, L., Deniz, O., Glomerulosclerosis identification in whole slide images using semantic segmentation, Computer methods and programs in biomedicine 184, 105273 (2020).
- [4] Cai, H., Li, J., Hu, M., Gan, C., Han, S., Efficientvit: Lightweight multi-scale attention for high-resolution dense prediction, Proceedings of the IEEE/CVF International Conference on Computer Vision, pg. 17302-17313. (2023).
- [5] Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A., Emerging properties in self-supervised vision transformers, Proceedings of the IEEE/CVF international conference on computer vision, pg. 9650-9660. (2021).
- [6] Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H., Encoder-decoder with atrous separable convolution for semantic image segmentation, arXiv:1802.02611 (2018).
- [7] Chen, R.J., Chen, C., Li, Y., Chen, T.Y., Trister, A.D., Krishnan, R.G., Mahmood, F., Scaling vision transformers to gigapixel images via hierarchical self-supervised learning, Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pg. 16144-16155. (2022).
- [8] Chen, R.J., Krishnan, R.G., Self-supervised vision transformers learn visual concepts in histopathology, arXiv preprint arXiv:2203.00585 (2022).
- [9] Chen, R.J., Lu, M.Y., Weng, W.H., Chen, T.Y., Williamson, D.F.,

- Manz, T., Shady, M., Mahmood, F., Multimodal co-attention transformer for survival prediction in gigapixel whole slide images, Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 4015-4025. (2021).
- [10] 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 (2023).
- [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, arXiv preprint arXiv:2402.19286 (2024).
- [12] Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al., An image is worth 16x16 words: Transformers for image recognition at scale, International Conference on Learning Representations (2020).
- [13] Esser, P., Rombach, R., Ommer, B., Taming transformers for high-resolution image synthesis, Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 12873-12883. (2021).
- [14] Feng, C., Liu, F., Artificial intelligence in renal pathology: Current status and future, Biomolecules and Biomedicine 23(2), 225 (2023).
- [15] Gadermayr, M., Dombrowski, A.K., Klinkhammer, B.M., Boor, P., Merhof, D., Cnn cascades for segmenting whole slide images of the kidney, arXiv preprint arXiv: 1708.00251 (2017).
- [16] Gu, A., Dao, T., Mamba: Linear-time sequence modeling with selective state spaces, arXiv preprint arXiv:2312.00752 (2023).
- [17] Guo, J., Hao, Z., Wang, C., Tang, Y., Wu, H., Hu, H., Han, K., Xu, C., Data-efficient large vision models through sequential autoregression, arXiv preprint arXiv:2402.04841 (2024).
- [18] Hara, S., Haneda, E., Kawakami, M., Morita, K., Nishioka, R., Zoshima, T., Kometani, M., Yoneda, T., Kawano, M., Karashima, S., et al., Evaluating tubulointerstitial compartments in renal biopsy specimens using a deep learning-based approach for classifying normal and abnormal tubules, PLoS one 17(7), e0271161 (2022).
- [19] Hatamizadeh, A., Tang, Y., Nath, V., Yang, D., Myronenko, A., Landman, B., Roth, H., Xu, D., Unetr: Transformers for 3d medical image segmentation, arXiv preprint arXiv:2103.10504 (2021).
- [20] He, Y., Nath, V., Yang, D., Tang, Y., Myronenko, A., Xu, D., Swinunetr-v2: Stronger swin transformers with stagewise convolutions for 3d medical image segmentation, International Conference on Medical Image Computing and Computer-Assisted Intervention, pp. 416-426. (2023).
- [21] Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H., nnu-net: a self-configuring method for deep learning-based biomedical image segmentation, Nature Methods 18(2), 203-211 (2021).
- [22] Israel, U., Marks, M., Dilip, R., Li, Q., Schwartz, M.S., Pradhan, E., Pao, E., Li, S., Pearson-Goulart, A., Perona, P., et al., A foundation model for cell segmentation, bioRxiv (2023).
- [23] Katharopoulos, A., Vyas, A., Pappas, N., Fleuret, F., Transformers are rns: Fast autoregressive transformers with linear attention, International conference on machine learning, pp. 5156-5165. (2020).
- [24] Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al., Segment anything, arXiv preprint arXiv:2304.02643 (2023).
- [25] Kumar, N., Verma, R., Anand, D., Zhou, Y., Onder, O.F., Tsougenis, E., Chen, H., Heng, P.A., Li, J., Hu, Z., et al., A multi-organ nucleus segmentation challenge, IEEE transactions on medical imaging 39(5), 1380-1391 (2019).
- [26] Ma, J., He, Y., Li, F., Han, L., You, C., Wang, B., Segment anything in medical images, Nature Communications 15(1), 654 (2024).
- [27] Ma, J., Xie, R., Ayyadhury, S., Ge, C., Gupta, A., Gupta, R., Gu, S., Zhang, Y., Lee, G., Kim, J., et al., The multi-modality cell segmentation challenge: Towards universal solutions, arXiv preprint arXiv:2308.05864 (2023).
- [28] Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al., Dinov2: Learning robust visual features without supervision, arXiv preprint arXiv:2304.07193 (2023).
- [29] Ronneberger, O., Fischer, P., Brox, T., U-Net: Convolutional Networks for Biomedical Image Segmentation, Proc. MICCAI (2015).
- [30] Shen, T., Zhang, Y., Qi, L., Kuen, J., Xie, X., Wu, J., Lin, Z., Jia, J., High quality segmentation for ultra high-resolution images, Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 1310-1319. (2022).
- [31] Tang, Y., Yang, D., Li, W., Roth, H.R., Landman, B., Xu, D., Nath, V., Hatamizadeh, A., Self-supervised pre-training of swin transformers for 3d medical image analysis, Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 20730-20740. (2022).
- [32] Van Den Oord, A., Vinyals, O., et al., Neural discrete representation learning, Advances in neural information processing systems 30 (2017).
- [33] Wang, J., Liu, Z., Zhao, L., Wu, Z., Ma, C., Yu, S., Dai, H., Yang, Q., Liu, Y., Zhang, S., et al., Review of large vision models and visual prompt engineering, Meta-Radiology, 100047 (2023).
- [34] Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P., Segformer: Simple and efficient design for semantic segmentation with transformers, Advances in Neural Information Processing Systems 34, 12077-12090 (2021).

Author Biography

Yucheng Tang received his Bachelor of Science from Tianjin University (2016), his Master of Science from New York University (2018), and his PhD from Vanderbilt University. He is currently a Sr. Research Scientist at NVIDIA. His work focuses on computationally efficient healthcare AI, with research areas including multi-modal foundation models, medical Vision-Language Models (VLMs), and their clinical translation.

JOIN US AT THE NEXT EI!

electronic IMAGING

Imaging across applications . . . Where industry and academia meet!



- **SHORT COURSES • EXHIBITS • DEMONSTRATION SESSION • PLENARY TALKS •**
- **INTERACTIVE PAPER SESSION • SPECIAL EVENTS • TECHNICAL SESSIONS •**

www.electronicimaging.org

