

StainNet: Scaling Self-Supervised Foundation Models on Immunohistochemistry and Special Stains for Computational Pathology

Jiawen Li^{*,1}

JW-LI24@MAILS.TSINGHUA.EDU.CN

Jiali Hu^{*,2}

72522269@CITYU-DG.EDU.CN

Xitong Ling^{*,1}

LINGXT23@MAILS.TSINGHUA.EDU.CN

Yongqiang Lv⁴

18146676064@163.COM

Yuxuan Chen¹

CHENYX23@MAILS.TSINGHUA.EDU.CN

Yizhi Wang¹

2022214454WANG@GMAIL.COM

Tian Guan^{†,1}

GUANTIAN@SZ.TSINGHUA.EDU.CN

Yifei Liu^{†,3}

78909008@QQ.COM

Yonghong He^{†,1}

HEYH@SZ.TSINGHUA.EDU.CN

¹ Tsinghua Shenzhen International Graduate School, Tsinghua University, China

² Biomedical Engineering Programme, City University of Hong Kong (Dongguan), China

³ Department of Pathology, Affiliated Hospital of Nantong University, Nantong, Jiangsu, China

⁴ Nantong University, Nantong, Jiangsu, China

Editors: Accepted for publication at MIDL 2026

Abstract

Foundation models trained with self-supervised learning (SSL) on large-scale histological images have significantly accelerated the development of computational pathology. These models can serve as backbones for region-of-interest (ROI) image analysis or patch-level feature extractors in whole-slide images (WSIs) based on multiple instance learning (MIL). Existing pathology foundation models (PFMs) are typically pre-trained on Hematoxylin-Eosin (H&E) stained pathology images. However, images such as immunohistochemistry (IHC) and special stains are also frequently used in clinical practice. PFMs pre-trained mainly on H&E-stained images may be limited in clinical applications involving these non-H&E images. To address this issue, we propose StainNet, a collection of self-supervised foundation models specifically trained for IHC and special stains in pathology images based on the vision transformer (ViT) architecture. StainNet contains a ViT-Small and a ViT-Base model, both of which are trained using a self-distillation SSL approach on over 1.4 million patch images extracted from 20,231 publicly available IHC and special staining WSIs in the HISTAI database. To evaluate StainNet models, we conduct experiments on three in-house slide-level IHC classification tasks, three in-house ROI-level special stain and two public ROI-level IHC classification tasks to demonstrate their strong ability. We also perform ablation studies such as few-ratio learning and retrieval evaluations, and compare StainNet models with recent larger PFMs to further highlight their strengths. The StainNet model weights are available at <https://github.com/WonderLandxD/StainNet>.

Keywords: Non-H&E images, foundation model, computational pathology.

* Contributed equally

† Corresponding author

1. Introduction

Modern computational pathology involves the use of deep learning models, such as convolutional neural networks (He et al., 2016) and ViT (Dosovitskiy, 2020), for clinical tasks on histopathological images, including cancer screening (Bejnordi et al., 2017), tumor analysis (Lu et al., 2021), morphological retrieval (Chen et al., 2022a), and survival prognosis (Chen et al., 2021a). In recent years, SSL methods have gained attention and have been proven to be highly effective (Kang et al., 2023). These methods design pretext tasks, such as contrastive learning (Chen et al., 2020) and image reconstruction (He et al., 2022), allowing the model to learn effective representations from unlabeled data and obtain good initialization weights, which can then be directly used as foundation models for pathology image feature extraction.

Most existing PFMs are pre-trained on H&E-stained histological images, which are routinely used and easily accessible in clinical workflows. As a result, these PFMs typically perform well with H&E-stained images. For example, by directly analyzing expert-annotated H&E-stained ROIs, PFMs can be used to identify tumor malignancy or microenvironment patterns (Lin et al., 2025). Additionally, when H&E resection or biopsy WSIs are cropped into patches, PFMs can extract these into high-dimensional features, which are used in weakly-supervised learning tasks based on MIL (Xu et al., 2025), such as cancer screening or biomarker detection (Campanella et al., 2025). However, there are other staining in clinical practice. For example, by using antibodies that bind to proteins in tissue cells, IHC slides are widely used to visualize the presence, location, and expression levels of specific proteins, which are commonly utilized for precision cancer diagnosis, disease subtyping, and drug selection. Due to significant differences in nuclear color, tissue patterns, and texture compared to H&E staining, directly transferring existing PFMs to these non-H&E images for representation may be limited in clinical performance (Li et al., 2024b).

To address this issue, we propose StainNet, a collection of PFMs specifically pre-trained on IHC and special stain images. Specifically, to train StainNet, we first select 20,231 non-H&E WSIs from the publicly available large-scale HISTAI database (Nechaev et al., 2025). For each WSI, we randomly crop up to 100 tissue-containing patch images to create our pre-training dataset. Next, to develop StainNet, we utilize DINO (Caron et al., 2021), a lightweight self-supervised learning framework specifically designed for ViT models, enabling StainNet to learn task-agnostic morphological information from a large collection of unlabeled special staining images. Finally, to evaluate StainNet, we construct downstream task datasets at both the ROI and WSI levels and compare the performance with existing pathology foundation models. Through this work, we aim to explore the adaptability of PFMs in clinical tasks involving non-H&E images and call for the broader exploration of computational pathology in more diverse clinical pathology image datasets.

2. Related works

2.1. Self-supervised PFMs

Foundation models are accelerating the advancement of computational pathology. By performing SSL on large-scale unlabeled pathology images, models can obtain histology-specific weights, which significantly improve their performance on downstream tasks. For instance,

CTransPath (Wang et al., 2022) uses a MoCov3-based SSL approach (Chen et al., 2021b) and is trained on 15.6 million image patches from 25 anatomical sites to demonstrate its ability in H&E staining image retrieval. HIPT (Chen et al., 2022b) and Lunit (Kang et al., 2023) adopt DINO (Caron et al., 2021), validating the feasibility of pre-training ViT models on histological images. Given the easy scalability of ViT parameters, recent works have aimed to train on larger datasets with larger ViT models. Examples include UNI (Chen et al., 2024), Virchow (Vorontsov et al., 2024), and Prov-Gigapath (Xu et al., 2024), which aim to diversify the pre-training dataset with more varied pathology images and leverage more computational resources to build stronger general pathology representation models. Additionally, some studies explore SSL in conjunction with special staining, such as Patho-Duet (Hua et al., 2024), which utilizes registered H&E and IHC images to help the model learn cross-staining features, aiding the transfer to more advanced nuclear information.

2.2. IHC and Special staining in clinical pathology

H&E staining is the most commonly used histological staining method, providing basic structural information by staining the cell nuclei and cytoplasm. However, H&E staining does not fully reveal certain cell features or tissue details, especially in the detection of immune responses or protein expression. To address these limitations, several staining methods such as IHC, Masson’s trichrome, and acid phosphatase staining are developed to highlight the specific proteins, cell subtypes, and extracellular matrix components (Gridley, 1957). For example, IHC utilizes antibodies to bind to target proteins, making it a valuable tool for cancer diagnosis and disease subtyping (Magaki et al., 2018); Masson’s trichrome staining is employed to visualize elastic fibers in tissues, particularly in fibrosis-related research (Lefkowitz, 2006). Recently, the development of computational pathology has accelerated research into the application of computer vision models to these special staining images, such as lymphocyte detection in IHC images (Swiderska-Chadaj et al., 2019) and the use of multiple special stains for kidney biopsy evaluation (Jayapandian et al., 2021). Some studies have also explored the direct use of generative models to convert H&E images into these stains to improve diagnostic accuracy (De Haan et al., 2021; Bai et al., 2023; Pati et al., 2024; Yan et al., 2023). However, these non-H&E images differ significantly from H&E staining in terms of staining mechanisms, texture features, and visual appearance, making it difficult for large-scale pre-trained models based on general H&E staining images to capture the specific features of special staining images (Li et al., 2024b).

3. Methodology

We develop and evaluate our StainNet models through data collection, model pre-training, and downstream testing. The entire process is illustrated in Figure 1.

3.1. Data collection

We use the publicly available large-scale WSI database HISTAI (Nechaev et al., 2025) as our pre-training data. The original HISTAI dataset contains 112,801 WSIs, of which 20,265 are non-H&E stains. We first apply the OTSU thresholding algorithm (Otsu et al., 1975) to exclude 34 special staining WSIs that can not extract foreground tissue, resulting in

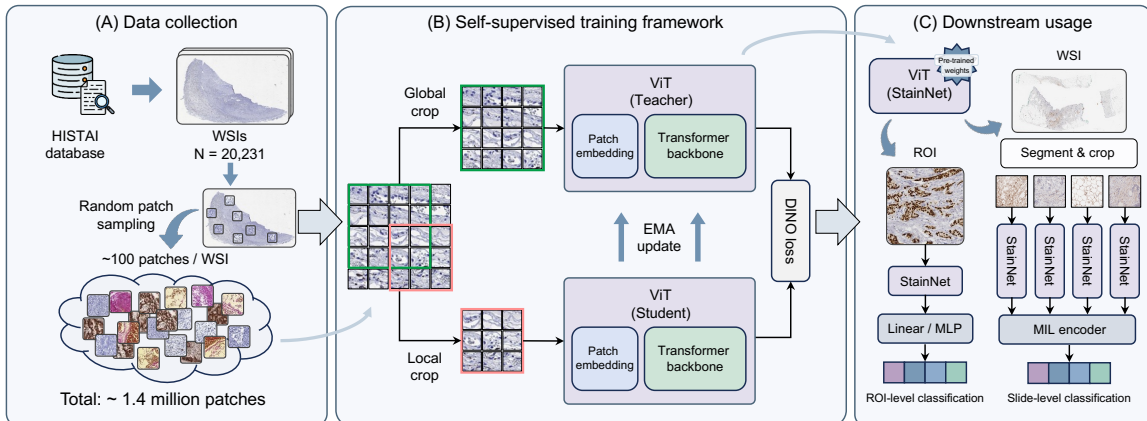


Figure 1: Overview of our proposed StainNet. We first collect approximately 1.4 million patch images with a resolution of 224×224 pixels from the HISTAI database. Then, we train StainNet models using the DINO SSL method. Finally, we adapt StainNet to downstream ROI and WSI tasks to explore its performance.

20,231 WSIs. Then, for each WSI, we perform non-overlapping 224×224 pixel sliding window operations on the foreground tissue, and randomly crop 100 image patches (if a WSI contains fewer than 100 patches, all the patches will be selected). In total, we obtain 1,418,938 patches for subsequent StainNet pre-training. More pre-training data details are provided in Appendix A.

3.2. Model pre-training

Our proposed StainNet models adopt the SSL method DINO (Caron et al., 2021) as the pre-training strategy to let the ViT model learn robust feature representations from IHC and special stain pathology images without manual annotations. DINO consists of two networks with the same architecture: a student network optimized via gradient descent and a teacher network updated using an exponential moving average (EMA) of the student weights. We follow the original DINO approach for data augmentation and multi-crop strategies, which are provided in Appendix B. Specifically, we generate a set of global and local views for each pathology image, and the student takes all views as input, while the teacher only receives the global views. By minimizing the cross-entropy loss between the output probability distributions of the two networks, the student is encouraged to infer global tissue context solely from local information. In addition, centering and sharpening operations are applied to the teacher outputs to prevent mode collapse. The detailed hyperparameters are provided in the implementation and Appendix B.

3.3. Downstream adaptation

Our pre-trained StainNet can be applied to both ROI and WSI analysis for special staining images. For ROI tasks, the extracted features from StainNet can be passed to either a single linear layer (linear probe) or a multilayer perceptron (MLP probe) consisting of two linear

layers with a $\text{ReLU}(\cdot)$ activation function for downstream classification. For WSI tasks, we first crop a set of tissue-containing patch images and then perform offline feature extraction using StainNet. The resulting $N \times C$ feature sequence (where N is the number of image patches and C is the feature dimension) can be fed into a series of MIL models to obtain slide-level representations and classification scores.

4. Experiments

4.1. Downstream datasets and implementation

We evaluate StainNet on three private IHC WSI tasks, NTUH-*Ki67*-Liver, *P53*-UCEC, and *MLH1*-UCEC, two public available IHC ROI tasks, BCI (Liu et al., 2022) and MIST (Li et al., 2023), and three in-house special stain ROI tasks, Glomerulus-Masson, Glomerulus-PAS, and Glomerulus-PASM, which are derived as subsets from our previous work (He et al., 2025). The NTUH-*Ki67*-Liver 3-class dataset consists of 46 *Ki67* benign liver slides, 43 *Ki67* hepatocellular carcinoma slides, and 50 *Ki67* normal liver slides. The *P53*-UCEC 5-class dataset consists of 139 *P53* wild-type, 38 *P53* mutant, 10 *P53* null, 14 *P53* subclonal, and 17 *P53* endometrial carcinoma slides. The *MLH1*-UCEC 3-class dataset consists of 30 *MLH1* deficient, 155 *MLH1* proficient, and 10 slides with partial loss of *MLH1* expression from endometrial carcinoma. All slides are collected from the Affiliated Hospital of Nantong University and scanned using the SQS-120P scanning system from Shenzhen Shengqiang Technology Co., Ltd. The annotations are performed and verified at the slide level by two pathologists. The BCI dataset is a collection of 4,870 paired H&E-IHC images showing *HER2* biomarker expression. We use the IHC images and categorize them into four *HER2* expression levels: 0, 1+, 2+, and 3+, with 240, 1,153, 2,142, and 1,335 samples, respectively. The MIST dataset is a large-scale H&E-IHC paired image dataset. We use the IHC images corresponding to four biomarkers: 5,642 *HER2*, 5,361 *Ki67*, 5,153 *ER*, and 5,139 *PR* to construct a 4-class dataset for evaluating the ability of PFMs to identify biomarkers across different staining conditions. The Glomerulus-Masson dataset consists of 200 normal, 57 slight, 112 moderate, and 113 severe Masson-stained glomeruli images. The Glomerulus-PAS dataset consists of 1200 normal, 1129 slight, 479 moderate, and 379 severe PAS-stained glomeruli images. The Glomerulus-PASM dataset consists of 200 normal, 76 slight, 135 moderate, and 87 severe PASM-stained glomeruli images.

We further evaluate StainNet models on two public H&E-stained ROI tasks, CRC-100K (Kather et al., 2019a) and KatherMS (Kather et al., 2019b), as well as one WSI task, PANDA-*Karo*. (Bulten et al., 2022), to investigate its capability on H&E image analysis. CRC-100K consists of nine categories: adipose, background, debris, lymphocytes, mucus, smooth muscle, normal colon mucosa, cancer-associated stroma, and colorectal adenocarcinoma epithelium. We perform label-stratified splitting on the official NCT-CRC-HE-100K dataset with a ratio of 0.8:0.2 to construct the train-val set, and use CRC-VAL-HE-7K as the test set. KatherMS contains two categories: microsatellite stable and microsatellite instable. We similarly conduct label-stratified splitting on the official training set with a ratio of 0.8:0.2 to obtain the train-val set, and use the official test set for evaluation. PANDA-*Karo*. consists of 5455 biopsy WSIs from Karolinska Institute, including 1924 G0, 1814 G1, 668 G2, 317 G3, 481 G4, and 251 G5. We label-stratify the train-val-test fold with ISUP grading ground truth of 0.5:0.2:0.3 for experiments.

Patch encoder <i>NTUH-Ki67-Liver</i>	MIL encoder						Overall
	ABMIL	SiMLP	TransMIL	WiKG	AMDMIL	S4MIL	
ResNet-50 (He et al., 2016)	0.742 _{0.093}	0.712 _{0.055}	0.829 _{0.055}	0.836 _{0.059}	0.838 _{0.071}	0.832 _{0.076}	0.7980
CTransPath (Wang et al., 2022)	0.668 _{0.050}	0.757 _{0.105}	0.858 _{0.040}	0.860 _{0.041}	0.802 _{0.018}	0.819 _{0.043}	0.7940
PathoDuet (Hua et al., 2024)	0.655 _{0.033}	0.641 _{0.024}	0.699 _{0.033}	0.738 _{0.032}	0.716 _{0.062}	0.725 _{0.030}	0.6957
HIPT (Chen et al., 2022b)	0.706 _{0.086}	0.753 _{0.053}	0.814 _{0.050}	0.815 _{0.058}	0.835 _{0.020}	0.843 _{0.017}	0.7943
Lunit (Kang et al., 2023)	0.747 _{0.048}	0.838 _{0.041}	0.872 _{0.051}	0.891 _{0.044}	0.872 _{0.020}	0.862 _{0.032}	0.8470
UNI (Chen et al., 2024)	0.764 _{0.055}	<u>0.891</u> _{0.052}	0.886 _{0.036}	0.914 _{0.033}	0.908 _{0.040}	0.890 _{0.052}	0.8755
PathOrchestra (Yan et al., 2025)	0.689 _{0.076}	0.837 _{0.068}	0.850 _{0.040}	0.861 _{0.093}	0.905 _{0.040}	0.878 _{0.054}	0.8367
GPFM (Ma et al., 2025)	0.754 _{0.060}	0.849 _{0.055}	0.899 _{0.070}	0.900 _{0.051}	0.900 _{0.049}	0.876 _{0.034}	0.8630
UNI-2 (Chen et al., 2024)	0.798 _{0.049}	0.842 _{0.087}	<u>0.907</u> _{0.041}	0.922 _{0.017}	0.922 _{0.046}	<u>0.909</u> _{0.033}	0.8833
Virchow-2 (Zimmermann et al., 2024)	0.814 _{0.050}	0.884 _{0.071}	0.895 _{0.067}	0.915 _{0.056}	0.907 _{0.022}	0.922 _{0.063}	<u>0.8895</u>
Prov-GigaPath (Xu et al., 2024)	0.852 _{0.098}	0.876 _{0.080}	0.884 _{0.077}	0.878 _{0.054}	<u>0.913</u> _{0.030}	0.905 _{0.047}	0.8847
H-optimus-0 (Saillard et al., 2024)	0.779 _{0.098}	0.841 _{0.057}	0.899 _{0.069}	0.859 _{0.047}	0.889 _{0.046}	0.888 _{0.045}	0.8592
StainNet-Small (Ours)	0.801 _{0.092}	0.873 _{0.042}	0.896 _{0.073}	0.916 _{0.030}	0.892 _{0.051}	0.883 _{0.060}	0.8768
StainNet-Base (Ours)	<u>0.830</u> _{0.084}	0.892 _{0.025}	0.928 _{0.050}	<u>0.921</u> _{0.048}	0.887 _{0.028}	0.900 _{0.048}	0.8930

Table 1: Comparison of balanced accuracy results for our StainNet with 12 PFMs on the NTUH-*Ki67*-Liver classification dataset. We use 6 different MIL methods for slide-level fine-tuning and report the average evaluation metrics with standard deviation across all folds. The best result is in **bold** and the second best is underlined.

For DINO training, we use the AdamW optimizer with a learning rate of 5×10^{-4} to pre-train StainNet for 100 epochs with a batch size of 256. For multi-cropping, global crops are sampled using a random resize scale in the range of 0.4-1.0, and 8 local crops are generated with a scale in the range of 0.05-0.4. For EMA design, we increase teacher momentum linearly from 0.9995 to 1.0 throughout training. The weight decay is scheduled from 0.04 to 0.4, and the teacher temperature is set to 0.04 without warm-up and last layer freezing. Our StainNet model family includes StainNet-Small and StainNet-Base, which are initialized with ImageNet-pre-trained (Deng et al., 2009) ViT-Small and ViT-Base backbones, and are then further pre-trained using DINO. We also enable normalization of the last layer to stabilize the training process. The StainNet-Small experiment is conducted on a single NVIDIA RTX PRO 6000 96GB GPU with mixed-precision and takes approximately 3.04 days to complete. The StainNet-Base experiment is conducted on 8 NVIDIA A100 80GB GPU and takes approximately 2.78 days to complete. For both StainNet-Small and StainNet-Base, we select the last epoch checkpoint as the final model weights.

For downstream fine-tuning, we use AdamW with a learning rate of 10^{-4} and a weight decay of 10^{-4} . We set 20 epochs with a batch size of 1 for the WSI task and 20 epochs with a batch size of 128 for the ROI task. For classification tasks, we save the fine-tuned model weights based on the best balanced accuracy, and also report accuracy, AUC, and F1-score from each fold. Early stopping is applied with a patience of 5 epochs, meaning that training is terminated if the balanced accuracy on the validation set does not improve for five consecutive epochs. For retrieval tasks, we perform image matching using cosine similarity and report Recall@K and mAP@K, where K ranges from 1 to 20. For NTUH-*Ki67*-Liver, *P53*-UCEC, *MLH1*-UCEC, BCI, and MIST, we adopt a label-stratified 5-fold

Patch encoder	<i>P53</i> -UCEC				<i>MLH1</i> -UCEC			
	Acc.	Bal acc.	AUC	F1-score	Acc.	Bal acc.	AUC	F1-score
ResNet-50 (He et al., 2016)	0.795 _{0.000}	0.333 _{0.000}	0.602 _{0.147}	0.704 _{0.000}	0.638 _{0.008}	0.200 _{0.000}	0.618 _{0.057}	0.497 _{0.011}
CTransPath (Wang et al., 2022)	0.795 _{0.000}	0.333 _{0.000}	0.705 _{0.061}	0.704 _{0.000}	0.785 _{0.024}	0.376 _{0.019}	0.779 _{0.094}	0.705 _{0.026}
PathoDuet (Hua et al., 2024)	0.795 _{0.000}	0.333 _{0.000}	0.362 _{0.153}	0.704 _{0.000}	0.762 _{0.041}	0.361 _{0.029}	0.689 _{0.087}	0.688 _{0.035}
HIPT (Chen et al., 2022b)	0.795 _{0.000}	0.333 _{0.000}	0.671 _{0.101}	0.704 _{0.000}	0.706 _{0.069}	0.285 _{0.080}	0.592 _{0.185}	0.600 _{0.102}
Lunit (Kang et al., 2023)	0.877 _{0.056}	0.538 _{0.117}	0.740 _{0.133}	0.840 _{0.081}	0.775 _{0.018}	0.369 _{0.026}	0.745 _{0.044}	0.695 _{0.014}
UNI (Chen et al., 2024)	<u>0.933</u> _{0.023}	<u>0.642</u> _{0.049}	0.833 _{0.064}	<u>0.908</u> _{0.024}	0.803 _{0.012}	0.393 _{0.012}	<u>0.853</u> _{0.024}	0.722 _{0.017}
PathOrchestra (Yan et al., 2025)	0.913 _{0.014}	0.616 _{0.039}	0.852 _{0.042}	0.889 _{0.015}	0.761 _{0.027}	0.367 _{0.046}	0.757 _{0.057}	0.684 _{0.034}
GPFM (Ma et al., 2025)	0.933 _{0.014}	0.642 _{0.029}	0.842 _{0.081}	0.909 _{0.014}	0.803 _{0.012}	0.393 _{0.012}	0.830 _{0.047}	0.722 _{0.016}
UNI-2 (Chen et al., 2024)	0.903 _{0.053}	0.567 _{0.114}	0.835 _{0.139}	0.870 _{0.063}	0.780 _{0.019}	0.379 _{0.020}	0.790 _{0.063}	0.700 _{0.015}
Virchow-2 (Zimmermann et al., 2024)	0.913 _{0.034}	0.616 _{0.049}	<u>0.862</u> _{0.040}	0.888 _{0.034}	0.775 _{0.023}	0.382 _{0.014}	0.815 _{0.054}	0.700 _{0.016}
Prov-GigaPath (Xu et al., 2024)	0.908 _{0.029}	0.614 _{0.047}	0.826 _{0.084}	0.884 _{0.028}	0.785 _{0.024}	<u>0.427</u> _{0.076}	0.850 _{0.047}	<u>0.724</u> _{0.033}
H-optimus-0 (Saillard et al., 2024)	0.903 _{0.033}	0.603 _{0.058}	0.804 _{0.110}	0.878 _{0.035}	0.785 _{0.011}	0.384 _{0.013}	0.808 _{0.064}	0.706 _{0.012}
StainNet-Small (Ours)	0.892 _{0.059}	0.562 _{0.135}	0.706 _{0.218}	0.854 _{0.087}	0.785 _{0.024}	0.373 _{0.029}	0.723 _{0.032}	0.704 _{0.023}
StainNet-Base (Ours)	0.933 _{0.014}	0.642 _{0.029}	0.880 _{0.068}	0.909 _{0.014}	<u>0.803</u> _{0.025}	0.536 _{0.120}	0.874 _{0.064}	0.754 _{0.034}

Table 2: Comparison results for our StainNet with 12 PFMs on the *P53*-UCEC and *MLH1*-UCEC classification datasets. We use ABMIL methods for slide-level fine-tuning and report four average evaluation metrics with standard deviation across all folds. The best result is in **bold** and the second best is underlined.

train-val cross-validation setting for evaluation. For Glomerulus-Masson, Glomerulus-PAS, and Glomerulus-PASM, we follow our previous work (Li et al., 2024b) to employ a label-stratified train-val-test split of 0.5:0.2:0.3 with five random seeds. For H&E tasks, we use the train-val-test setting with three random seeds. We use a single NVIDIA RTX 3090 24GB GPU with full-precision computing to run all downstream tasks. For MIL and PFM baselines, we follow their official implementations, including hyperparameter settings and data preprocessing pipelines. Except for these model-specific configurations, all other downstream fine-tuning hyperparameters are kept identical to those used in StainNet.

4.2. Comparison results on WSI tasks

For the WSI task, we compare the proposed StainNet-Small and StainNet-Base with 12 PFMs from small to large parameter size: (1) ResNet-50 (ImageNet pre-trained) (He et al., 2016), (2) CTransPath (Wang et al., 2022), (3) PathoDuet (Hua et al., 2024), (4) HIPT (Chen et al., 2022b), (5) Lunit (Kang et al., 2023), (6) UNI (Chen et al., 2024), (7) PathOrchestra (Yan et al., 2025), (8) GPFM (Ma et al., 2025), (9) UNI-2 (Chen et al., 2024), (10) Virchow-2 (Zimmermann et al., 2024), (11) Prov-GigaPath (Xu et al., 2024), and (12) H-optimus-0 (Saillard et al., 2024). For the WSI task, we evaluate six MIL models: (1) ABMIL (Ilse et al., 2018), a gated-attention MIL, (2) SiMLP (Li et al., 2025), a method combining unsupervised average pooling with an MLP classifier, (3) TransMIL (Shao et al., 2021), a transformer-based MIL, (4) WiKG (Li et al., 2024a), a dynamic graph representation method, (5) AMDMIL (Ling et al., 2024), a method combining masked denoising with agent tokens, and (6) S4MIL (Fillioux et al., 2023), a state space model-based MIL. Table 1 shows the comparison results of the StainNet series models on the six MIL models on NTUH-*Ki67*-Liver.

Patch encoder	Linear probe				MLP probe			
	<i>Acc.</i>	<i>Bal acc.</i>	<i>AUC</i>	<i>F1-score</i>	<i>Acc.</i>	<i>Bal acc.</i>	<i>AUC</i>	<i>F1-score</i>
<i>BCI</i>								
ResNet-50 (He et al., 2016)	0.587 _{0.010}	0.407 _{0.010}	0.774 _{0.018}	0.541 _{0.014}	0.709 _{0.016}	0.614 _{0.019}	0.900 _{0.008}	0.703 _{0.016}
CTransPath (Wang et al., 2022)	0.575 _{0.011}	0.387 _{0.007}	0.808 _{0.007}	0.499 _{0.008}	0.736 _{0.015}	0.653 _{0.021}	0.913 _{0.008}	0.733 _{0.015}
PathoDuet (Hua et al., 2024)	0.612 _{0.016}	0.430 _{0.015}	0.810 _{0.014}	0.554 _{0.028}	0.687 _{0.019}	0.564 _{0.044}	0.877 _{0.012}	0.664 _{0.024}
HIPT (Chen et al., 2022b)	0.564 _{0.023}	0.404 _{0.023}	0.754 _{0.020}	0.532 _{0.028}	0.661 _{0.009}	0.533 _{0.010}	0.859 _{0.010}	0.650 _{0.011}
Lunit (Kang et al., 2023)	0.696 _{0.015}	0.545 _{0.018}	0.880 _{0.008}	0.681 _{0.017}	0.872 _{0.016}	0.816 _{0.033}	0.976 _{0.002}	0.871 _{0.017}
UNI (Chen et al., 2024)	0.785 _{0.020}	0.693 _{0.031}	0.936 _{0.011}	0.781 _{0.021}	0.900 _{0.011}	0.877 _{0.024}	0.984 _{0.004}	0.900 _{0.011}
PathOrchestra (Yan et al., 2025)	0.732 _{0.021}	0.601 _{0.024}	0.914 _{0.009}	0.721 _{0.022}	0.847 _{0.012}	0.804 _{0.035}	0.969 _{0.004}	0.846 _{0.012}
GPFM (Ma et al., 2025)	0.704 _{0.004}	0.528 _{0.005}	0.902 _{0.008}	0.681 _{0.004}	0.844 _{0.017}	0.816 _{0.026}	0.966 _{0.006}	0.845 _{0.017}
UNI-2 (Chen et al., 2024)	0.772 _{0.010}	0.624 _{0.019}	0.934 _{0.007}	0.759 _{0.012}	0.874 _{0.013}	0.852 _{0.019}	0.979 _{0.003}	0.874 _{0.013}
Virchow-2 (Zimmermann et al., 2024)	0.822 _{0.013}	<u>0.735</u> _{0.018}	0.962 _{0.001}	0.828 _{0.007}	<u>0.913</u> _{0.009}	<u>0.888</u> _{0.011}	0.990 _{0.002}	<u>0.913</u> _{0.009}
Prov-GigaPath (Xu et al., 2024)	0.784 _{0.018}	0.673 _{0.022}	0.941 _{0.008}	0.778 _{0.019}	0.885 _{0.009}	0.865 _{0.016}	0.981 _{0.002}	0.884 _{0.009}
H-optimus-0 (Saillard et al., 2024)	0.789 _{0.031}	0.650 _{0.028}	0.941 _{0.012}	0.779 _{0.031}	0.878 _{0.010}	0.838 _{0.023}	0.978 _{0.003}	0.876 _{0.011}
StainNet-Small (Ours)	0.709 _{0.025}	0.597 _{0.036}	0.892 _{0.023}	0.701 _{0.025}	0.874 _{0.010}	0.839 _{0.020}	0.979 _{0.004}	0.873 _{0.010}
StainNet-Base (Ours)	<u>0.811</u> _{0.006}	0.744 _{0.007}	<u>0.951</u> _{0.003}	<u>0.809</u> _{0.005}	0.913 _{0.008}	0.896 _{0.014}	<u>0.989</u> _{0.001}	0.913 _{0.008}
Patch encoder								
<i>MIST</i>								
ResNet-50 (He et al., 2016)	0.597 _{0.009}	0.592 _{0.009}	0.835 _{0.002}	0.593 _{0.008}	0.671 _{0.005}	0.667 _{0.004}	0.884 _{0.002}	0.671 _{0.005}
CTransPath (Wang et al., 2022)	0.663 _{0.008}	0.658 _{0.008}	0.878 _{0.004}	0.659 _{0.006}	0.732 _{0.007}	0.728 _{0.007}	0.920 _{0.004}	0.731 _{0.007}
PathoDuet (Hua et al., 2024)	0.590 _{0.010}	0.585 _{0.011}	0.827 _{0.006}	0.586 _{0.013}	0.664 _{0.013}	0.660 _{0.012}	0.885 _{0.004}	0.664 _{0.011}
HIPT (Chen et al., 2022b)	0.587 _{0.007}	0.581 _{0.007}	0.830 _{0.006}	0.577 _{0.007}	0.681 _{0.006}	0.676 _{0.007}	0.896 _{0.004}	0.676 _{0.010}
Lunit (Kang et al., 2023)	0.738 _{0.006}	0.733 _{0.006}	0.922 _{0.003}	0.736 _{0.005}	0.871 _{0.005}	0.869 _{0.006}	0.977 _{0.001}	0.871 _{0.006}
UNI (Chen et al., 2024)	0.773 _{0.007}	0.769 _{0.007}	0.940 _{0.002}	0.771 _{0.008}	0.874 _{0.005}	0.872 _{0.005}	0.979 _{0.001}	0.874 _{0.005}
PathOrchestra (Yan et al., 2025)	0.760 _{0.008}	0.755 _{0.009}	0.934 _{0.003}	0.758 _{0.009}	0.846 _{0.004}	0.843 _{0.005}	0.969 _{0.001}	0.845 _{0.005}
GPFM (Ma et al., 2025)	0.730 _{0.009}	0.725 _{0.009}	0.916 _{0.003}	0.726 _{0.009}	0.812 _{0.006}	0.808 _{0.006}	0.956 _{0.002}	0.811 _{0.006}
UNI-2 (Chen et al., 2024)	0.782 _{0.006}	0.778 _{0.007}	0.943 _{0.002}	0.782 _{0.007}	0.842 _{0.002}	0.838 _{0.002}	0.967 _{0.001}	0.841 _{0.003}
Virchow-2 (Zimmermann et al., 2024)	<u>0.820</u> _{0.005}	<u>0.816</u> _{0.005}	<u>0.959</u> _{0.002}	<u>0.820</u> _{0.005}	<u>0.918</u> _{0.006}	<u>0.917</u> _{0.006}	<u>0.989</u> _{0.001}	<u>0.919</u> _{0.006}
Prov-GigaPath (Xu et al., 2024)	0.803 _{0.004}	0.799 _{0.004}	0.951 _{0.003}	0.802 _{0.004}	0.873 _{0.004}	0.871 _{0.005}	0.978 _{0.002}	0.873 _{0.005}
H-optimus-0 (Saillard et al., 2024)	0.752 _{0.007}	0.748 _{0.007}	0.930 _{0.001}	0.753 _{0.008}	0.841 _{0.003}	0.838 _{0.003}	0.966 _{0.002}	0.841 _{0.003}
StainNet-Small (Ours)	0.779 _{0.006}	0.775 _{0.006}	0.941 _{0.002}	0.778 _{0.007}	0.884 _{0.006}	0.882 _{0.006}	0.982 _{0.002}	0.884 _{0.006}
StainNet-Base (Ours)	0.825 _{0.005}	0.822 _{0.005}	0.960 _{0.002}	0.825 _{0.006}	0.923 _{0.005}	0.921 _{0.005}	0.991 _{0.001}	0.923 _{0.005}

Table 3: Comparison results for our StainNet with 12 PFMs on the BCI and MIST classification datasets. We use linear probe and MLP probe methods for ROI-level fine-tuning and report four average evaluation metrics with standard deviation across all folds. The best result is in **bold** and the second best is underlined.

Table 2 shows the comparison results of ABMIL on *P53*-UCEC and *MLH1*-UCEC. We observe that the StainNet-Base achieves overall best performance on the three WSI tasks. We also report the paired t-test p-values between StainNet-Base and other PFMs on the *P53*-UCEC dataset, which are shown in Appendix C. For example, when compared with billion-scale models such as Prov-GigaPath and H-optimus-0, the 86M-parameter StainNet-Base still outperforms them on NTUH-Ki67-Liver by 0.83% and 3.38%. Moreover, it outperforms them by 10.9% and 15.2% in balanced accuracy on *MLH1*-UCEC. It is noteworthy that although StainNet-Small has extremely small parameters (only 22M), it remains highly competitive. For instance, on NTUH-Ki67-Liver, it achieves overall best performance compared to PFMs with similar parameter size, and outperforms UNI and GPFM (303M) by 0.13% and 1.38%. These results demonstrate the advantages of large-scale pre-training on non-H&E pathology images.

Patch encoder	<u>Glomerulus-Masson</u>		<u>Glomerulus-PAS</u>		<u>Glomerulus-PASM</u>	
	<i>Bal acc.</i>	<i>F1-score</i>	<i>Bal acc.</i>	<i>F1-score</i>	<i>Bal acc.</i>	<i>F1-score</i>
ResNet-50 (He et al., 2016)	0.253 _{0.005}	0.254 _{0.024}	0.336 _{0.004}	0.416 _{0.004}	0.258 _{0.011}	0.271 _{0.013}
CTransPath (Wang et al., 2022)	0.284 _{0.022}	0.292 _{0.059}	0.478 _{0.010}	0.551 _{0.006}	0.273 _{0.027}	0.273 _{0.053}
PathoDuet (Hua et al., 2024)	0.296 _{0.028}	0.315 _{0.038}	0.353 _{0.016}	0.431 _{0.018}	0.256 _{0.019}	0.277 _{0.035}
HIPT (Chen et al., 2022b)	0.319 _{0.048}	0.279 _{0.062}	0.528 _{0.017}	0.574 _{0.012}	0.296 _{0.066}	0.300 _{0.076}
Lunit (Kang et al., 2023)	0.380 _{0.087}	0.396 _{0.163}	0.607 _{0.008}	0.648 _{0.009}	0.361 _{0.044}	0.377 _{0.132}
UNI (Chen et al., 2024)	<u>0.478</u> _{0.037}	<u>0.537</u> _{0.040}	0.632 _{0.030}	<u>0.665</u> _{0.023}	0.390 _{0.036}	0.446 _{0.033}
UNI-2 (Chen et al., 2024)	0.395 _{0.015}	0.441 _{0.015}	0.535 _{0.012}	0.577 _{0.012}	0.301 _{0.053}	0.328 _{0.074}
Virchow-2 (Zimmermann et al., 2024)	0.359 _{0.025}	0.418 _{0.030}	0.555 _{0.006}	0.611 _{0.006}	0.403 _{0.057}	0.456 _{0.067}
Prov-GigaPath (Xu et al., 2024)	0.430 _{0.058}	0.484 _{0.066}	0.625 _{0.019}	0.661 _{0.012}	0.501 _{0.025}	0.559 _{0.023}
H-optimus-0 (Saillard et al., 2024)	0.427 _{0.042}	0.480 _{0.052}	<u>0.633</u> _{0.014}	0.660 _{0.010}	0.427 _{0.042}	0.480 _{0.043}
StainNet-Small (Ours)	0.351 _{0.026}	0.347 _{0.099}	0.592 _{0.020}	0.631 _{0.014}	0.335 _{0.065}	0.368 _{0.106}
StainNet-Base (Ours)	0.505 _{0.027}	0.567 _{0.028}	0.669 _{0.020}	0.693 _{0.017}	<u>0.475</u> _{0.063}	<u>0.533</u> _{0.058}

Table 4: Comparison of balanced accuracy and F1-score results for our StainNet with 10 PFMs on the Glomerulus-Masson, Glomerulus-PAS, and Glomerulus-PASM classification dataset. We use linear probe for ROI-level fine-tuning and report the average evaluation metrics with standard deviation across all random seeds. The best result is in **bold** and the second best is underlined.

4.3. Comparison results on ROI tasks

Table 3 and Table 4 present the comparative results on the IHC tasks and the special stain glomerulus tasks. For the IHC tasks, we observe that Virchow-2, which is pre-trained on both H&E and IHC images, and StainNet-Base both achieve superior overall performance compared to other models. However, considering that Virchow-2 is based on a ViT-Huge/14 architecture, its computational cost is substantially higher than that of StainNet-Base. Meanwhile, StainNet-Small also demonstrates a strong balance between computational efficiency and performance when compared with models of similar or even larger scales. It is worth noting that our goal is not to directly replace previous PFMs, but rather to highlight the advantages of pre-training on domain-specific pathology images. We further believe that training on larger-scale non-H&E datasets and using models with bigger parameter size will lead to further performance gains, which we leave as an important direction for future work. For the special stain tasks, we find that the advantage of StainNet-Small diminishes when compared with models of similar scale; however, StainNet-Base further widens the performance gap with other larger PFMs. We also provide Umap visualization for BCI in Appendix D.

4.4. Comparison of image retrieval

We also examine the retrieval capability of StainNet-Small and StainNet-Base to assess their semantic representation quality on special staining histology images. We compare StainNet with CTransPath and UNI on the MIST dataset, as shown in Figure 2. We observe that both StainNet models consistently achieve the best Recall@K and mAP@K

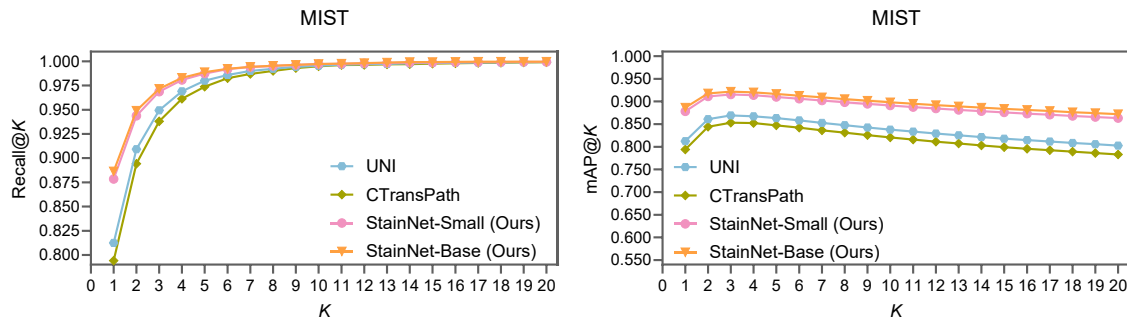


Figure 2: Comparison of image retrieval performance among CTransPath, UNI, and our StainNet-Small and StainNet-Base on the MIST ROI-level dataset.




Figure 3: Linear probe and MIL probe fine-tuning results of CTransPath, UNI and our StainNet on the MIST ROI-level dataset across varying training data ratios.

across all choices of K , indicating that they have learned a well-structured embedding space for special staining images.

4.5. Effectiveness of different fine-tuning data ratios

We further evaluate the fine-tuning performance of StainNet models under varying amounts of downstream special-stain training data to assess their learning capability. Figure 3 presents the results of StainNet-Small, StainNet-Base, CTransPath, and UNI on the MIST dataset using linear and MLP probes under different proportions of fine-tuning samples. We observe that StainNet-Base achieves the best performance across all ratio settings. Although StainNet-Small performs slightly worse than UNI under the 10% data setting with a linear probe, it achieves the best performance in all other settings and substantially outperforms the CTransPath of similar parameter size. Notably, applying an MLP probe significantly increases the performance gap between StainNet and UNI in low-fine-tuning data scenarios.

Datasets	CRC-100K	KatherMS	PANDA- <i>Karo</i> .
HIPT (Chen et al., 2022b)	0.895 _{0.001}	0.554 _{0.010}	0.217 _{0.005}
CTransPath (Wang et al., 2022)	0.937 _{0.001}	0.688 _{0.002}	0.238 _{0.009}
Lunit (Kang et al., 2023)	0.910 _{0.003}	0.694 _{0.003}	0.270 _{0.011}
UNI (Chen et al., 2024)	0.937 _{0.001}	0.711 _{0.004}	0.535 _{0.014}
StainNet-Small (Ours)	0.931 _{0.005}	0.684 _{0.001}	0.252 _{0.005}
StainNet-Base (Ours)	0.930 _{0.004}	0.658 _{0.006}	0.322 _{0.009}

Table 5: Comparison of balanced accuracy for our StainNet with 4 PFMs on two H&E-stained ROI-level and one WSI-level classification dataset. We use linear probe for ROI-level and ABMIL for WSI-level fine-tuning, and report the average evaluation metrics with standard deviation across all random seeds. The best result is in **bold** and the second best is underlined.

4.6. Comparison results on H&E-stained tasks

We further explore the potential of the proposed StainNet models on H&E image tasks. Table 5 presents the quantitative results. We find that StainNet-Small is not significantly different from other PFMs with similar parameter sizes on both tasks, and that StainNet-Base does not improve performance on H&E image tasks despite having a larger parameter size than StainNet-Small. However, the large PFM UNI, specifically trained on H&E-stained images, leads in both task types, especially on the WSI task, where the gap is significant.

4.7. Comparison of different pre-training data ratios and iterations

We also conduct an ablation study to investigate the impact of different pre-training data scales. Specifically, we randomly sample 25 and 50 patches from each WSI to construct two pre-training datasets of different sizes, which are then used to pretrain StainNet-Small with DINO. Then we adapt the pre-trained models to the BCI and MIST ROI-level downstream tasks, as shown in Figure 4.a. We observe that larger pre-training datasets lead to stronger downstream performance, which validates the scalability of DINO on non-H&E images.

Finally, we investigate the performance variation of StainNet-Small under different pre-training iterations and the behavior of its loss curve. We believe these observations can provide valuable insights for the development of SSL methods in computational pathology, where pre-training details such as the expected loss curve and convergence behavior are often underreported. Figure 4.b and Figure 4.c illustrate the pre-training loss curve over 100 epochs and the corresponding linear probe performance on the MIST and Glomerulus-PAS dataset at six iteration checkpoints. We observe that on MIST, the loss plateaus between epoch 30 and 50, but with significant performance improvement during this period. After epoch 60, the overall loss begins to decrease again, but with slight changes in performance. Ultimately, the model reaches its optimal state at the final iteration. On Glomerulus-PAS, performance increases progressively with the number of epochs. These findings suggest that the shift in the DINO pre-training loss may serve as a valuable indicator of representation quality. While additional exploration is required to generalize these observations to other settings, we encourage future SSL development in computational pathology to con-




Figure 4: Pre-training loss curve of StainNet and comparison of linear probe performance on MIST under six different iterations.

sider allocating more training epochs, which may be crucial for achieving optimal model performance.

5. Conclusion and future works

In this work, we propose StainNet, a collection of pathology foundation models specifically for special staining histopathology image representation. Utilizing the DINO self-supervised framework, we train StainNet-Small and StainNet-Base models based on over 1.4 million unlabeled images from non-H&E WSIs in HISTAI. We evaluate StainNet models on three in-house IHC WSI tasks, two public IHC ROI tasks, and three in-house special stain ROI tasks, including classification, few-ratio settings, and image retrieval. The results demonstrate that StainNet models learn strong and robust representations for IHC and special stain pathology images.

For future work, we plan to develop larger StainNet models by expanding the pre-training dataset, scaling the model parameter size, and evaluating them on broader downstream tasks. We will also explore integrating StainNet with H&E-based PFMs for multi-modal applications such as survival prognosis and therapy response prediction. We believe that StainNet will further accelerate the research and applications in computational pathology involving diverse histological modalities.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (NSFC) (82430062), the Shenzhen Engineering Research Centre (XMHT20230115004), and the Tsinghua Shenzhen International Graduate School Cross-disciplinary Research and Innovation Fund Research Plan (JC2024002). We thank Qiehe Sun for providing GPU support.

References

- Bijie Bai, Xilin Yang, Yuzhu Li, Yijie Zhang, Nir Pillar, and Aydogan Ozcan. Deep learning-enabled virtual histological staining of biological samples. *Light: Science & Applications*, 12(1):57, 2023.
- Babak Ehteshami Bejnordi, Mitko Veta, Paul Johannes Van Diest, Bram Van Ginneken, Nico Karssemeijer, Geert Litjens, Jeroen AWM Van Der Laak, Meyke Hermsen, Quirine F Manson, Maschenka Balkenhol, et al. Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer. *Jama*, 318(22): 2199–2210, 2017.
- Wouter Bulten, Kimmo Kartasalo, Po-Hsuan Cameron Chen, Peter Ström, Hans Pinckaers, Kunal Nagpal, Yuannan Cai, David F Steiner, Hester Van Boven, Robert Vink, et al. Artificial intelligence for diagnosis and gleason grading of prostate cancer: the panda challenge. *Nature medicine*, 28(1):154–163, 2022.
- Gabriele Campanella, Neeraj Kumar, Swaraj Nanda, Siddharth Singi, Eugene Fluder, Ricky Kwan, Silke Muehlstedt, Nicole Pfarr, Peter J Schüffler, Ida Häggström, et al. Real-world deployment of a fine-tuned pathology foundation model for lung cancer biomarker detection. *Nature Medicine*, 31(9):3002–3010, 2025.
- Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021.
- Chengkuan Chen, Ming Y Lu, Drew FK Williamson, Tiffany Y Chen, Andrew J Schaumberg, and Faisal Mahmood. Fast and scalable search of whole-slide images via self-supervised deep learning. *Nature Biomedical Engineering*, 6(12):1420–1434, 2022a.
- Richard J Chen, Ming Y Lu, Muhammad Shaban, Chengkuan Chen, Tiffany Y Chen, Drew FK Williamson, and Faisal Mahmood. Whole slide images are 2d point clouds: Context-aware survival prediction using patch-based graph convolutional networks. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 339–349. Springer, 2021a.
- Richard J Chen, Chengkuan Chen, Yicong Li, Tiffany Y Chen, Andrew D Trister, Rahul G Krishnan, and Faisal Mahmood. Scaling vision transformers to gigapixel images via hierarchical self-supervised learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16144–16155, 2022b.

- Richard J Chen, Tong Ding, Ming Y Lu, Drew FK Williamson, Guillaume Jaume, Andrew H Song, Bowen Chen, Andrew Zhang, Daniel Shao, Muhammad Shaban, et al. Towards a general-purpose foundation model for computational pathology. *Nature medicine*, 30(3): 850–862, 2024.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PmLR, 2020.
- Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9640–9649, 2021b.
- Kevin De Haan, Yijie Zhang, Jonathan E Zuckerman, Tairan Liu, Anthony E Sisk, Miguel FP Diaz, Kuang-Yu Jen, Alexander Nobori, Sofia Liou, Sarah Zhang, et al. Deep learning-based transformation of h&e stained tissues into special stains. *Nature communications*, 12(1):4884, 2021.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- Leo Fillioux, Joseph Boyd, Maria Vakalopoulou, Paul-Henry Cournède, and Stergios Christodoulidis. Structured state space models for multiple instance learning in digital pathology. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 594–604. Springer, 2023.
- Mary Francis Gridley. *Manual of histologic and special staining technics*. Armed Forces Institute of Pathology, 1957.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022.
- Qiming He, Yingming Xu, Qiang Huang, Yanxia Wang, Jing Ye, Yonghong He, Jing Li, Lianghai Zhu, Zhe Wang, and Tian Guan. Unveiling pathology-related predictive uncertainty of glomerular lesion recognition using prototype learning. *Journal of Biomedical Informatics*, 161:104745, 2025.
- Shengyi Hua, Fang Yan, Tianle Shen, Lei Ma, and Xiaofan Zhang. Pathoduet: Foundation models for pathological slide analysis of h&e and ihc stains. *Medical Image Analysis*, 97: 103289, 2024.

- Maximilian Ilse, Jakub Tomczak, and Max Welling. Attention-based deep multiple instance learning. In *International conference on machine learning*, pages 2127–2136. PMLR, 2018.
- Catherine P Jayapandian, Yijiang Chen, Andrew R Janowczyk, Matthew B Palmer, Clarissa A Cassol, Miroslav Sekulic, Jeffrey B Hodgin, Jarcy Zee, Stephen M Hewitt, John O’Toole, 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.
- Mingu Kang, Heon Song, Seonwook Park, Donggeun Yoo, and Sérgio Pereira. Benchmarking self-supervised learning on diverse pathology datasets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3344–3354, 2023.
- Jakob Nikolas Kather, Johannes Krisam, Pornpimol Charoentong, Tom Luedde, Esther Herpel, Cleo-Aron Weis, Timo Gaiser, Alexander Marx, Nektarios A Valous, Dyke Ferber, et al. Predicting survival from colorectal cancer histology slides using deep learning: A retrospective multicenter study. *PLoS medicine*, 16(1):e1002730, 2019a.
- Jakob Nikolas Kather, Alexander T Pearson, Niels Halama, Dirk Jäger, Jeremias Krause, Sven H Loosen, Alexander Marx, Peter Boor, Frank Tacke, Ulf Peter Neumann, et al. Deep learning can predict microsatellite instability directly from histology in gastrointestinal cancer. *Nature medicine*, 25(7):1054–1056, 2019b.
- Jay H Lefkowitz. Special stains in diagnostic liver pathology. In *Seminars in diagnostic pathology*, volume 23, pages 190–198. Elsevier, 2006.
- Fangda Li, Zhiqiang Hu, Wen Chen, and Avinash Kak. Adaptive supervised patchnce loss for learning h&e-to-ihc stain translation with inconsistent groundtruth image pairs. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 632–641. Springer, 2023.
- Jiawen Li, Yuxuan Chen, Hongbo Chu, Qiehe Sun, Tian Guan, Anjia Han, and Yonghong He. Dynamic graph representation with knowledge-aware attention for histopathology whole slide image analysis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11323–11332, 2024a.
- Jiawen Li, Tian Guan, Qingxin Xia, Yizhi Wang, Xitong Ling, Jing Li, Qiang Huang, Zihan Wang, Zhiyuan Shen, Yifei Ma, et al. Unlocking adaptive digital pathology through dynamic feature learning. *arXiv preprint arXiv:2412.20430*, 2024b.
- Jiawen Li, Jiali Hu, Qiehe Sun, Renao Yan, Minxi Ouyang, Tian Guan, Anjia Han, Chao He, and Yonghong He. Can we simplify slide-level fine-tuning of pathology foundation models? *arXiv preprint arXiv:2502.20823*, 2025.
- Yixuan Lin, Weiping Lin, Chenxu Guo, Xinxin Yang, Hongxue Meng, and Liansheng Wang. Tumor microenvironment-guided fine-tuning of pathology foundation models for esophageal squamous cell carcinoma immunotherapy response prediction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 660–669. Springer, 2025.

- Xitong Ling, Minxi Ouyang, Yizhi Wang, Xinrui Chen, Renao Yan, Hongbo Chu, Junru Cheng, Tian Guan, Sufang Tian, Xiaoping Liu, et al. Agent aggregator with mask denoise mechanism for histopathology whole slide image analysis. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 2795–2803, 2024.
- Shengjie Liu, Chuang Zhu, Feng Xu, Xinyu Jia, Zhongyue Shi, and Mulan Jin. Bci: Breast cancer immunohistochemical image generation through pyramid pix2pix. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1815–1824, 2022.
- Ming Y Lu, Tiffany Y Chen, Drew FK Williamson, Melissa Zhao, Maha Shady, Jana Lipkova, and Faisal Mahmood. Ai-based pathology predicts origins for cancers of unknown primary. *Nature*, 594(7861):106–110, 2021.
- Jiabo Ma, Zhengrui Guo, Fengtao Zhou, Yihui Wang, Yingxue Xu, Jinbang Li, Fang Yan, Yu Cai, Zhengjie Zhu, Cheng Jin, et al. A generalizable pathology foundation model using a unified knowledge distillation pretraining framework. *Nature Biomedical Engineering*, pages 1–20, 2025.
- Shino Magaki, Seyed A Hojat, Bowen Wei, Alexandra So, and William H Yong. An introduction to the performance of immunohistochemistry. *Biobanking: methods and protocols*, pages 289–298, 2018.
- Dmitry Nechaev, Alexey Pchel'nikov, and Ekaterina Ivanova. Histai: An open-source, large-scale whole slide image dataset for computational pathology. *arXiv preprint arXiv:2505.12120*, 2025.
- Nobuyuki Otsu et al. A threshold selection method from gray-level histograms. *Automatica*, 11(285-296):23–27, 1975.
- Pushpak Pati, Sofia Karkampouna, Francesco Bonollo, Eva Comp erat, Martina Radi c, Martin Spahn, Adriano Martinelli, Martin Wartenberg, Marianna Kruithof-de Julio, and Marianna Rapsomaniki. Accelerating histopathology workflows with generative ai-based virtually multiplexed tumour profiling. *Nature machine intelligence*, 6(9):1077–1093, 2024.
- Charlie Saillard, Rodolphe Jenatton, Felipe Llinares-L pez, Zelda Mariet, David Cahan , Eric Durand, and Jean-Philippe Vert. H-optimus-0, 2024. URL <https://github.com/bioptimus/releases/tree/main/models/h-optimus/v0>.
- Zhuchen Shao, Hao Bian, Yang Chen, Yifeng Wang, Jian Zhang, Xiangyang Ji, et al. Transmil: Transformer based correlated multiple instance learning for whole slide image classification. *Advances in neural information processing systems*, 34:2136–2147, 2021.
- Zaneta Swiderska-Chadaaj, Hans Pinckaers, Mart Van Rijthoven, Maschenka Balkenhol, Margarita Melnikova, Oscar Geessink, Quirine Manson, Mark Sherman, Antonio Polonia, Jeremy Parry, et al. Learning to detect lymphocytes in immunohistochemistry with deep learning. *Medical image analysis*, 58:101547, 2019.

- Eugene Vorontsov, Alican Bozkurt, Adam Casson, George Shaikovski, Michal Zelechowski, Kristen Severson, Eric Zimmermann, James Hall, Neil Tenenholtz, Nicolo Fusi, et al. A foundation model for clinical-grade computational pathology and rare cancers detection. *Nature medicine*, 30(10):2924–2935, 2024.
- Xiyue Wang, Sen Yang, Jun Zhang, Minghui Wang, Jing Zhang, Wei Yang, Junzhou Huang, and Xiao Han. Transformer-based unsupervised contrastive learning for histopathological image classification. *Medical image analysis*, 81:102559, 2022.
- Hanwen Xu, Naoto Usuyama, Jaspreet Bagga, Sheng Zhang, Rajesh Rao, Tristan Naumann, Cliff Wong, Zelalem Gero, Javier González, Yu Gu, et al. A whole-slide foundation model for digital pathology from real-world data. *Nature*, 630(8015):181–188, 2024.
- Hongming Xu, Mingkang Wang, Duanbo Shi, Huamin Qin, Yunpeng Zhang, Zaiyi Liu, Anant Madabhushi, Peng Gao, Fengyu Cong, and Cheng Lu. When multiple instance learning meets foundation models: advancing histological whole slide image analysis. *Medical Image Analysis*, 101:103456, 2025.
- Fang Yan, Jianfeng Wu, Jiawen Li, Wei Wang, Yirong Chen, Linda Wei, Jiaxuan Lu, Wen Chen, Zizhao Gao, Jianan Li, et al. Pathorchestra: A comprehensive foundation model for computational pathology with over 100 diverse clinical-grade tasks. *npj Digital Medicine*, 8(1):695, 2025.
- Rena Yan, Qiming He, Yiqing Liu, Peng Ye, Lianghai Zhu, Shanshan Shi, Jizhou Gou, Yonghong He, Tian Guan, and Guangde Zhou. Unpaired virtual histological staining using prior-guided generative adversarial networks. *Computerized Medical Imaging and Graphics*, 105:102185, 2023.
- Eric Zimmermann, Eugene Vorontsov, Julian Viret, Adam Casson, Michal Zelechowski, George Shaikovski, Neil Tenenholtz, James Hall, David Klimstra, Razik Yousfi, et al. Virchow2: Scaling self-supervised mixed magnification models in pathology. *arXiv preprint arXiv:2408.00738*, 2024.

Appendix A. Details of pre-training datasets

We observe that the non-H&E WSI data collected in HISTAI exhibit a certain degree of ambiguity in naming conventions. Specifically, heterogeneous clone identifiers and inconsistent file naming rules are used across different cases, and some WSIs lack explicit annotations of their corresponding non-H&E staining types. As a result, it is practically challenging to assign fully accurate and unified stain labels to all WSIs. Nevertheless, we made systematic efforts to curate and organize the data to the best extent possible, with the goal of providing the community with a reasonably accurate distribution for the pre-training data, thereby reflecting the overall composition and diversity of non-H&E staining modalities. Based on our statistics, the HISTAI non-H&E pre-training dataset comprises 331 distinct immunohistochemistry (IHC) staining types and 12 special staining types. Figure 5 illustrates the overall distribution of different staining modalities within the pre-training data. Furthermore, Table 6 reports the distribution of IHC staining types with more than 20 cases, highlighting the major IHC sources contributing to the pre-training process. Table 7 summarizes the case distribution of all special staining types. We hope that these statistics and visualizations will serve as a useful reference for future studies leveraging HISTAI non-H&E data.




Figure 5: Overview of the StainNet pre-training dataset. **a.** Distribution of the top 20 IHC stains by case count. **b.** Distribution of the all special stains by case count. **c.** Distribution of patch image types used for training.

STAINNET

IHC stain type	Case number	Patch image number
Ki67	1419	95715
ER	762	51533
PR	691	46757
CD20	535	39513
CD138	464	41006
HER2/neu	441	22850
MCK	431	28564
CD3 (VET)	403	29128
CD10	385	35373
CK7	365	25864
Bcl-2	313	28075
Bcl-6	311	27345
S100 (4c4)	289	24579
CD23	287	24359
CD34	282	23313
PD-L1	281	25500
Synaptophysin	272	18537
CD5	269	20015
P63-7JUL	253	14431
HER2	249	8607
Cyclin-D1	247	22386
MUM1-Protein	233	20339
TTF-1	230	13932
CK20 (ks20)	214	14266
CD30	208	17529
SMA	202	18070
PAX-8	190	14440
Chromogranin-A	185	13375
CD45	183	16063
GATA-3	178	11796
CDX2	174	10253
Desmin	154	10415
MSH6	154	10044
MIH1	151	11546
SOX-10	142	14610
MSH2	141	11224
PAX-5	138	10074
PMS2	136	10648
CD56	128	9803
Vimentin	126	11039
Melan-A	124	10128

IHC stain type	Case number	Patch image number
P53	123	9530
P16	119	7881
HMB45	116	9334
CD56 (123c3)	112	7871
P40	112	6112
AMACR	112	4823
Epstein-Barr-Birus	111	10141
CK14	104	5918
CD117	100	8255
WT1-WT49	100	7267
CD68	99	8079
E-Cadherin	92	4808
CK5&6	90	5444
S100	81	7301
CD15	79	6646
CD4	70	5575
Kappa-Light-Chain	67	5637
Lambda-Light-Chain	65	4974
EMA (gp1)	63	5318
CD8	61	4992
A1K	59	4542
DOG-1	58	4482
CD79A	56	4088
CD99	54	5074
Calponin-1	54	4412
Calretinin	52	4395
Arginase-1	49	4299
SATB2	49	3918
CD16	47	4400
ERG	47	3070
HIA-DR-Antigen	46	3525
Glypican-3	46	3104
CK8 & CK18 (b22)	46	2957
Napsin-a	46	2906
INSML	43	3268
CK-HMW	43	1532
GFAP	42	3566
P63	42	3167
CD31	40	3207
CD2	38	3012
Beta-Catenin	37	3231
ER-other	37	3037
CD21	36	2869

IHC stain type	Case number	Patch image number
INI-1	35	3383
CD7	35	2468
TDT	32	2934
CK19	32	2278
Androgen-Receptor	32	1747
Inhibin-Alpha	29	2646
Mammaglobin	29	2089
PSA	29	1796
C-MYC	27	2393
HHV8	27	2067
CD3-K	27	1636
Myeloperoxidase	26	2424
CD61	26	2096
CK20 (KS20)	26	2071
STAT6	26	1992
MDM2	26	1850
CD38	26	1643
Heppar	25	2213
PLAP	25	1968
LIFR	23	2249
IGG4	23	2106
Caix-TH22	23	2059
CK-5-6-d5-16b4	23	1653
CEA	22	1677
Alpha-Fetoprotein	21	1850
SAll4	21	1504
CD3	21	1453
Podoplanin	20	1845
Total	16150	1216942

Table 6: IHC stain type distribution of our data for StainNet pre-training.

Special stain type	Case number	Patch image number
Giemsa	869	114016
Argentum	67	2956
Van Gieson	38	5039
PAS	30	3061
Alcian Blue	31	3743
Congo Red	17	763
Toluidine Blue	16	1011
Calcitonin	9	816
Ziehl Neelsen	8	881
PAS-D	1	100
Romanovsky	1	100
Grocott	1	14
Total	1088	132500

Table 7: Special stain type distribution of our data for StainNet pre-training.

Appendix B. Details of DINO pre-training settings

We provide the data augmentation strategy and detailed hyperparameter settings for DINO, as shown in Table 8 and Table 9.

Augmentation	Value
Image size	224×224
Local view size	96×96
Global random resize scale	(0.4, 1.0)
Local random resize scale	(0.05, 0.4)
Number of local views	8
Random horizontal flip	0.5
Random vertical flip	0.0
Random gray scale	0.2
Color jitter brightness	0.8
Color jitter contrast	0.8
Color jitter hue	0.2
Color jitter prob	0.8
Color jitter saturation	0.4
Color jitter strength	0.5
Gaussian blur limit	0
Gaussian blur prob	1.0
Gaussian blur sigmas	(0.1, 2.0)
Global view Gaussian blur prob	0.1
Global view Solarize prob	0.2
Global view Solarize threshold	0.5
Local view: Gaussian blur prob	0.5
Normalization mean	(0.5, 0.5, 0.5)
Normalization std	(0.5, 0.5, 0.5)

Table 8: DINO data augmentation settings used in StainNet pre-training.

We provided some examples of original images and their augmented versions, which are illustrated in Figure 6.




Figure 6: A set of original special staining images and their data-augmented versions.

Hyperparameter	Value
Batch normalization	False
Bottleneck dimension	256
Center momentum	0.9
Hidden dimension	2048
Learning rate scale method	linear
Momentum end	1.0
Momentum start	0.9995
Last layer normalization	True
Output dimension	65536
Epoch	100
Total batch size	256
Student freeze last layer epochs	0
Student temperature	0.1
Teacher temperature	0.04
Warmup teacher temperature	0.04
Warmup teacher temperature epochs	0
Weight decay end	0.4
Weight decay start	0.04
Optimizer	AdamW
Optimizer betas	(0.9, 0.999)
Optimizer eps	1e-8
Optimizer learning rate	5e-4
Optimizer weight decay	0.04
Automaticmixedprecision	BF16

Table 9: DINO hyperparamters used in StainNet pre-training. StainNet-Samll is trained using $1 \times 96\text{GB}$ NVIDIA PRO 9000, and StainNet-Base is trained using $8 \times 80\text{GB}$ NVIDIA A100.

Appendix C. Statistical significance analysis

Model	Acc. (p)	Bal acc. (p)	AUC (p)	F1-score (p)
ResNet-50	2.70e-04***	3.34e-03**	1.16e-03**	1.03e-04***
CTransPath	2.40e-01	3.13e-02*	1.36e-01	3.85e-02*
PathoDuet	8.26e-03**	2.57e-02*	2.37e-03**	5.93e-03**
HIPT	4.66e-02*	2.48e-02*	3.08e-02*	4.43e-02*
Lunit	3.27e-02*	2.85e-02*	8.04e-04***	7.86e-03**
UNI	9.94e-01	6.03e-02	4.93e-01	1.40e-01
PathOrchestra	3.68e-02*	5.65e-02	2.23e-02*	3.61e-02*
GPFM	9.94e-01	6.03e-02	1.77e-01	1.44e-01
UNI-2	8.96e-02	3.40e-02*	7.33e-03**	6.78e-03**
Virchow-2	3.33e-02*	3.89e-02*	6.68e-02	3.13e-03**
Prov-GigaPath	1.61e-02*	1.12e-01	4.99e-01	6.55e-02
H-optimus-0	1.78e-01	4.66e-02*	9.12e-02	4.50e-02*
StainNet-Small	9.83e-02	2.60e-02*	2.10e-03**	1.72e-02*

Table 10: Paired t-test p-values between StainNet-Base and different PFMs on the *P53*-UCEC dataset (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$).

Appendix D. Feature visualization

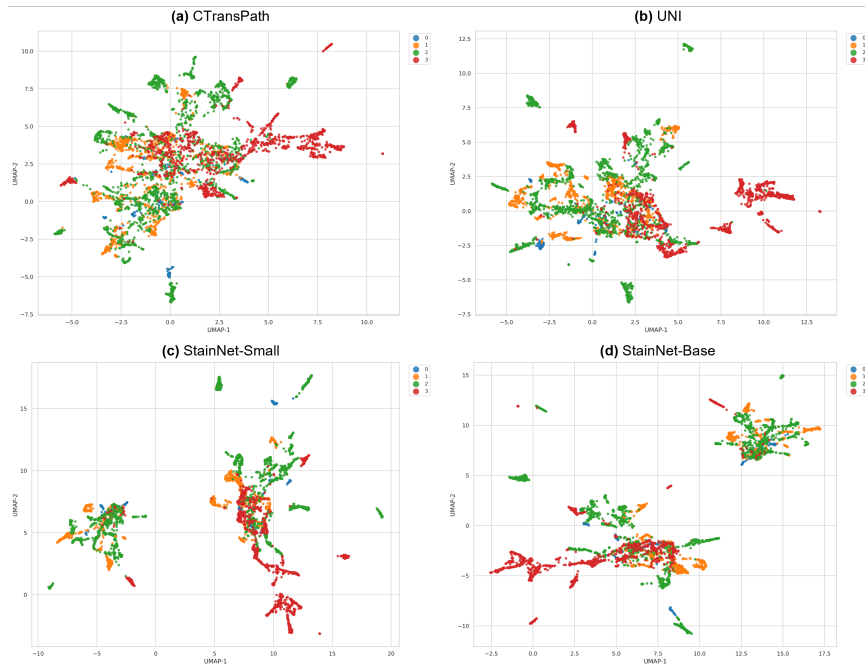


Figure 7: Umap visualization of high-dimension feature