

Vim4Path: Self-Supervised Vision Mamba for Histopathology Images

Ali Nasiri-Sarvi¹ Vincent Quoc-Huy Trinh² Hassan Rivaz³ Mahdi S. Hosseini¹

¹Department of Computer Science and Software Engineering (CSSE), Concordia University, Canada

²Institute for Research in Immunology and Cancer of the University of Montreal, Canada

³Department of Electrical and Computer Engineering (ECE), Concordia University, Canada

Abstract

Representation learning from Gigapixel Whole Slide Images (WSI) poses a significant challenge in computational pathology due to the complicated nature of tissue structures and the scarcity of labeled data. Multi-instance learning methods have addressed this challenge, leveraging image patches to classify slides utilizing pretrained models using Self-Supervised Learning (SSL) approaches. The performance of both SSL and MIL methods relies on the architecture of the feature encoder. This paper proposes leveraging the Vision Mamba (Vim) architecture, inspired by state space models, within the DINO framework for representation learning in computational pathology. We evaluate the performance of Vim against Vision Transformers (ViT) on the Camelyon16 dataset for both patch-level and slide-level classification. Our findings highlight Vim's enhanced performance compared to ViT, particularly at smaller scales, where Vim achieves an 8.21 increase in ROC AUC for models of similar size. An explainability analysis further highlights Vim's capabilities, which reveals that Vim uniquely emulates the pathologist workflow—unlike ViT. This alignment with human expert analysis highlights Vim's potential in practical diagnostic settings and contributes significantly to developing effective representation-learning algorithms in computational pathology. We release the codes and pretrained weights at <https://github.com/AtlasAnalyticsLab/Vim4Path>.

1. Introduction

Representation learning from Gigapixel Whole Slide Image (WSI) is a challenging problem in computational pathology. Each tissue WSI comprises hundreds of thousands of cells with biological importance, interacting with each other as neighborhoods, with microscopic juxtaposition of different cell types. To represent the holistic nature of each WSI, one needs to collect thousands of manageable instances (known as image patches) from each slide to process and interpret the

tissue cells in WSI. However, the diagnostic relevance of the tissue cells within each image patch is unidentified, where only a weakly-supervised label is available on the slide-level for representation. To bridge the gap, Multi-Instance Learning (MIL) methods [11] are adopted to aggregate embedded features from image patches and classify them on the bag level (i.e., slide) [25]. Most MIL methods rely on utilizing a feature extractor, such as ResNet50 [22], as a backbone model to assign an embedding to each item in the bag. Pre-trained models from natural images, such as ImageNet [10], were traditionally adopted for feature extraction. However, the distribution of pathology images differs from natural images, causing sub-optimal performance when using pre-trained ImageNet weights.

To address these issues, better approaches exist beyond simply fine-tuning ImageNet weights, as WSIs often use weak supervision with only one label per slide, and typically, just a few image patches contain diagnostic information. Although more granular annotations can be used (e.g., region-of-interest (ROI) or patch-level), they are costly and scarce due to the extensive expert analysis they require [25]. As a result, recent advancements in computational pathology have increasingly concentrated on Self-Supervised Learning (SSL) approaches [28]. These methods train on extensive volumes of unlabeled data to achieve robust feature representation without relying on labels. The features obtained can then be downstream to tasks like classification, utilizing a much smaller set of labeled data. Additionally, it is important to recognize that the choice of the encoder architecture significantly influences the effectiveness of the embeddings produced by SSL methods.

Recently, a new architecture design inspired by state space models, known as Vision Mamba (Vim) [48], has emerged. In Section 4, we will discuss why this architecture holds promise for pathology applications. It combines the advantageous inductive biases of CNNs, requiring less scale and computational power, with the ability to perceive long-range dependencies, similar to ViTs. Our research intends to explore and adapt the Vim architecture within SSL schemes. To evaluate its performance, we employ DINO [3], a well-

known SSL framework used for ViT models, as a benchmark to contrast the capabilities of Vim against ViT architectures. Our study is divided into four main objectives and findings:

1. We compare the two architectures in both patch-level and slide-level classification tasks on the Camelyon16 [1] dataset for benchmarking.
2. We show that the Vim model significantly outperforms the ViT model at a smaller scale.
3. We demonstrate that the Vim model is competitive and often surpasses the ViT model’s performance at scale.
4. We perform explainability testing and discover that the Vim model mimics the pathologist’s method of diagnosing different image parts.

The rest of paper reads as follows. Section 2 discusses the related works. Section 3 discusses preliminaries from state space modeling. In Section 4, we explain our proposed method and the intuition behind Vim models for pathology datasets. Finally, in Section 5, we empirically compare the Vim and ViT architectures, assess their performance at different scales, and perform explainability analysis.

2. Related Work

Whole-slide images (WSIs) in pathology datasets are Gigapixel images, with individual samples reaching pixel sizes up to 150,000x150,000, hindering the direct application of deep learning techniques without adequate adjustments. To manage these Gigapixel images, they are broken into smaller patches, with limited field-of-view, for processing using Multi-Instance Learning (MIL) methods [11] such as [27], [32], [26], [38], [29], [47]. They mainly use a pretrained encoder from natural image datasets such as ImageNet [10] or Pathology datasets such as Camelyon16 [1]. These pre-trained networks on pathology datasets are usually trained in a self-supervised manner due to a lack of granular labels. The main idea behind self-supervised learning is to train an encoder without labels on a large data cohort to obtain a good feature representation, then transfer the weights for downstream tasking to train a classifier head on a smaller labeled dataset. Methods such as DINO [3], SimCLR [5, 6], MoCo [7, 8, 23], and MAE [24] are some of the self-supervised learning frameworks that were introduced for natural image processing. In computational pathology, contrastive learning is used to train self-supervised models such as [30], [41], [9]. HIPT [4] uses a hierarchical structure of ViT [12] models with each level in the hierarchy trained through DINO [3] and ViT models [12] architecture. A comprehensive review is provided in [25] on the utility of self-supervised learning in computational pathology applications.

Recently, an alternative for ViT [12] has been proposed based on state space models. Vision Mamba (Vim) [48] extends state space models [16–18, 20, 21, 35] to images, delivering state-of-the-art performance with linear memory requirements. While there have been previous works in using

state space models in pathology images such as [15], [13], and [46], the utility of the state space models as an image encoder within the realm of computational pathology has not been yet explored.

3. Preliminaries

In this section, we review some concepts regarding the state space models, Vision Mamba, and Self-Supervised Learning Methods. These preliminaries are needed to discuss our proposed method in Sec 4.

3.1. State Space Models

State space models are mathematical models for describing continuous linear systems. The state equation

$$h'(t) = Ah(t) + Bx(t) \quad (1)$$

captures how the state $h(t)$ changes overtime based on the input $x(t)$. The output equation

$$y(t) = Ch(t) + Dx(t) \quad (2)$$

links the output to the hidden state and input.

This model based on A , and B is for continuous input. However, in deep learning models, the inputs are discrete. To discretize the model, a time scale parameter Δ is used to get discrete parameters \bar{A} and \bar{B} followed by

$$\bar{A} = \exp(\Delta A), \quad \bar{B} = (\Delta A)^{-1}(\exp(\Delta A) - I)\Delta B \quad (3)$$

In the S4 model [19], the parameters A , B , C , and Δ are time-invariant, meaning they are defined independently of each individual token in the sequence. This design is mainly to prevent slow down of recurrent modeling, allowing this method to use a global convolution to model the sequential data. While this time-invariance helps with processing speed, it prevents the model from behaving based on the input, hence limiting its performance.

3.2. Vision Mamba

In the selective state space (Mamba) [17], the parameters B , C , and Δ are computed based on the input sequence using a linear projection, making them time-variant, leaving A the only time-invariant parameter. This prevents the usage of a global convolution, forcing this method to use recurrent processing, which does not utilize the parallelization of GPUs. To increase the processing speed, this method uses a hardware-aware algorithm to utilize GPUs for recurrent modeling. Mamba [17] is developed for 1D sequences such as text and audio. Vision Mamba (Vim) [48] extends this method to images by extracting small image tiles (for example, 16×16 tiles, which are also known as tokens), projecting them using a linear layer to an embedding dimension similar to [12]. Furthermore, to incorporate positional

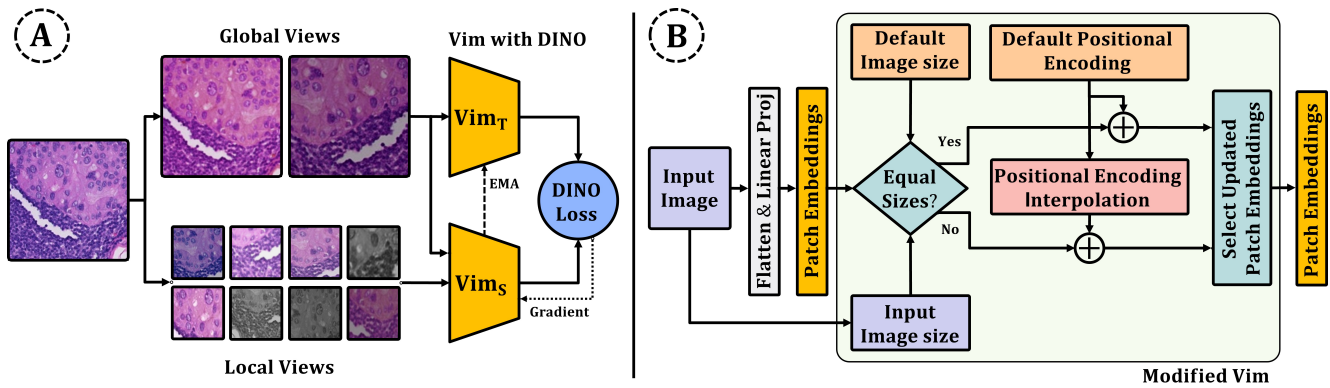


Figure 1. Detailed architecture of VIM within the DINO framework. We modify the Vim model to adapt to input image size for positional embedding interpolation and employ the modified model within DINO as a backbone architecture for self-supervised learning.

awareness, they add positional encoding to the extracted projections similar to [12]. Since Mamba [17] is a sequential model, it uses a bidirectional Mamba to better capture the representation for a given token based on future and previous tokens since we have access to all tokens in images.

3.3. Self-Supervised Learning using DINO

DINO [3] utilizes self-distillation in a teacher-student setup, where both networks share the same architecture but differ in parameters. The teacher’s outputs guide the student’s training without labeled data. Training involves cropping the input image into global and local views and applying augmentations to each view, similar to SimCLR. The student predicts the teacher’s output from all views, while the teacher only sees the global ones. Teacher’s parameters are updated as an exponential moving average of the student weights, ensuring stability and enhancing representation quality. This method’s emphasis on global-local view training is crucial for pathology images, as it helps in understanding the “global to local” relationships vital for detecting transitions in cancer states across neighboring cells, offering a more accurate representation of disease progression.

3.4. Multi-Instance Learning using CLAM

We use CLAM [32] as a framework to compare different architectures on slide-level classification. This method employs a data-efficient approach under a weakly supervised setting. It utilizes attention-based multiple-instance learning, allowing it to identify sub-regions within the slides most indicative of the slide-level label, effectively focusing on the most informative features without requiring detailed annotation. Our reason for using it is its ability to achieve high diagnostic accuracy with minimal data and labeling effort.

4. Methodology

This section elaborates on our proposed method for adopting a modified Vim [48] with Positional Encoding Interpolation [12] within the DINO [3] framework. We also explain why Vim [48] is a good alternative to ViT [12] architecture.

4.1. Vision Mamba within DINO

Since the DINO framework requires different image sizes (e.g., 224×224 and 96×96), any encoder in this framework should be able to receive different image sizes. This is not the case with the Vanilla Vim [48], and we introduce a Positional Encoding Interpolation module similar to [12] to the Vim [48] model. The positional encoding interpolation module works by dynamically resizing the positional encoding to match the dimensions of the input image, ensuring spatial information is accurately maintained across different scales. This is achieved by using bicubic interpolation to adjust the patch embeddings according to the calculated target grid size derived from the input dimensions and patch size. This adaptation allows the default Vim model to handle images of varying sizes, allowing us to adopt it within the DINO [3] framework. Detailed architecture of Vim within the DINO framework using the Positional Encoding Interpolation is shown in Fig 1.

4.2. Architecture Design Comparison

CNNs are designed with an inductive bias that processes adjacent areas within an image, such as segments A and B (Fig 2), using fixed-size kernels. However, these models struggle with long-range dependencies due to their limited receptive field. This constrains their capacity to directly relate distant areas (e.g. segment A with Z) depicted in Fig 2. In contrary, ViTs lack this inductive bias, treating every image patch the same, with only positional encoding differing to retain spatial relationships. This design allows ViTs to effectively capture long-range dependencies, such as the connection

between segments A and Z in Fig 2. However, this comes at the cost of requiring more data and computational resources to learn the nuances of local features, which are naturally handled by the inductive bias in CNNs.

Vim combines the strengths of both CNNs and ViTs. By employing sequential processing, Vim uses an inductive bias focusing on local dependencies between segments like A and B in Fig 2. Additionally, Vim’s usage of Mamba [17] overcomes the short context limitation inherent in previous image sequential models like PixelRNN [42], making it capable of capturing long-range dependencies across images, similar to ViT’s capacity to relate segment A with segment Z in Fig 2. This combination allows Vim to leverage the benefits of both localized and global information in images, as it sequentially scans through the image data while maintaining a long-range perspective, as demonstrated in Fig 2.




Figure 2. Comparison between different architecture designs. Vim sequential processing allows the model to capture both short-range and long-range dependencies.

4.3. Vim Relationship to Pathology

Sequential processing of Vision Mamba on patch level from the whole slide image is depicted in Fig 3. Each patch is divided into multiple tokens, and the Vision Mamba raster scans the tokens from each patch in sequence form following a lawnmower pattern. This scan enables the encoder to pick up short-range dependencies for representation in patch embedding. Later, the embedded features are aggregated on the slide level for MIL representation. This process imitates the realistic representation of cancer cells in sequence, from local zooming to the global representation on the slide, where a similar slide navigation is used by the pathologists for diagnosing cancer cells under the microscopy [33].

Pathology insights to short-range dependencies. Capturing both short- and long-range dependencies in pathology images is crucial. In fact, adding the sequences for short-range dependencies at higher magnification (e.g. 10x) is more important compared to low magnification (e.g. 5x). The most important event is the sequence of cells next to each other at higher magnification, as that is what a pathologist actually practice for clinical diagnosis. We empirically demonstrate this effect in our experiments in Section 5. While sequences at lower magnification are relevant, at

high magnification it is where it becomes essential. To elaborate this in detail, during cancer evolution, normal cells throughout the body incur an initial driver mutation into cancerous cells [36]. This transition does not occur in isolation within a single cell population as coordinated responses from neighboring cells shape the tissue into sequential patterns that interact with the mutation-burdened tumor cells [36]. These sequences of neighboring cells are at the core of modern clinical pathology and genomics research [2]. For instance, the sequence of lymphocytes adjacent to tumor cells and their expression of the protein PD-L1 underlie response to immune checkpoint inhibitors, a discovery based on the proximity of these immune cells to cancer, at the basis of the 2018 Nobel Prize in Physiology or Medicine [44]. Indeed, the most novel and costly form of genomics research, termed “spatial transcriptomics,” was developed to overcome the lack of tools that interrogate cell networks and cell-to-cell sequences on pathology tissues [34]. In a very short period since its inception, adding these spatial sequences has led to thousands of novel discoveries [34]. Overall, there is a medical and biological need for computational pathology to deploy algorithms that query neighboring events. While both ViT and Vim are capable of encoding long-range dependencies, applying short-range dependencies is key to unraveling the complex cell networks in pathology and medicine.

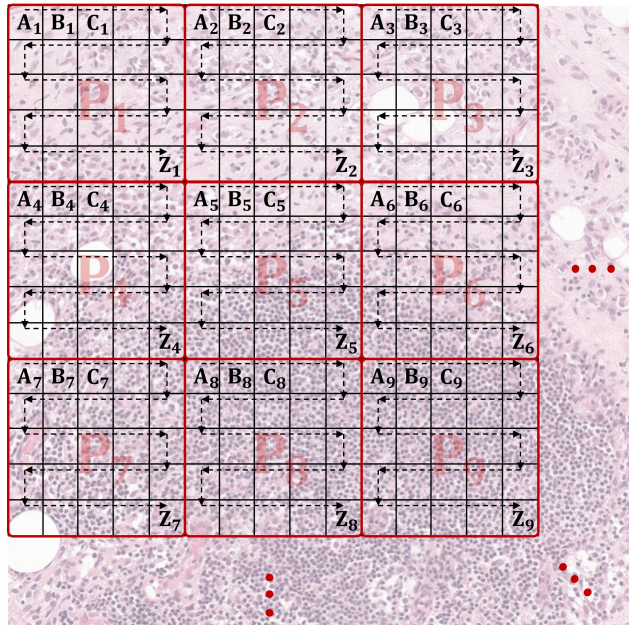


Figure 3. Sequential processing of Vim done on each patch level from slide for feature embedding. This is similar to the lawnmower pattern used for slide navigation by pathologists to study cellular neighbourhoods in the tissue for cancer diagnosis. The information from each patch (i.e. embeddings) are put together to reach to a consensus on the slide level (i.e. aggregation).

5. Experiments

In this section, we explain in detail the procedures and results of our experiments. Specifically, we outline the dataset used and the processing pipeline for this dataset, our experimental setup, and the results obtained from our experiments.

5.1. Dataset

We used Camelyon16 [1] dataset as a benchmark for our study, as extensively described in [1]. Table 1 summarizes the detailed distribution across different categories for the used dataset.

Category	Training Set	Testing Set
Normal Slides	158	80
Tumor Slides	110	49

Table 1. Distribution of the used Camelyon16 dataset

For the pre-training phase, patches are only extracted from the training set slides, with label information being ignored, using the pre-processing pipeline of CLAM [32] with the default hyper-parameters. Table 2 summarizes the distribution of the extracted patches for pre-training across two zooming levels. Note that the tumor patches are obtained from tumor slides without reference to ROI (region of interest) labels, indicating that not all extracted images necessarily represent tumor patches. The main reason for using 5x and 10x zooming levels was to process the data with limited computation power. We aim to compare ViT and Vim models in the context of pathology images using the same experimental setup. Hence, the 20x zooming level, which is the default for many pathology methods, was not explored in this study (due to computation limits).

Zooming Level	Normal Patches	Tumor Patches
5x	146,508	21,446
10x	555,231	446,608

Table 2. Distribution of the pretraining dataset used. Note that the tumor patches are extracted from tumor slides without considering ROI (region of interest) labels, hence each individual image is likely not a tumor patch.

To test each model’s patch classification performance, we perform similarly to the PCam dataset [43] to ensure we have a balanced patch classification dataset with labels. We do not use the original PCam [43] dataset since the image sizes are 96×96 . For the normal slides, we extract patches as before, using the pre-processing pipeline of CLAM [32]. We extract patches from ROI regions with 10% overlap for the tumor slides. We ensure that the extracted patches have a minimum of 50% intersection with the ROI region. Finally, we select a subset of patches from the normal patches to

ensure a balanced dataset by uniformly sampling from all the slides to ensure maximum slide diversity among the chosen patches. We further validate the labels of the extracted patches through an expert pathologist. For the 10x zooming level with a patch image size of 224, we call this dataset PCam-224-10x. The statistics of these datasets are available in Table 3.

Dataset	Training Patches	Testing Patches
PCam-224-5x	2×7602	2×6611
PCam-224-10x	2×28955	2×25370

Table 3. Distribution of the extracted patch classification datasets. We ensure to have the same number of patches in both classes (shown by $2 \times$)

Model	Parameters (Millions)
ViT-ti	6
ViT-s	22
Vim-ti	7
Vim-ti-plus	13
Vim-s	26

Table 4. Number of Trainable Parameters for each model (rounded to the nearest Million)

5.2. Experimental setup

Our experiments were conducted using 4 NVIDIA V100 GPUs, with a batch size of 128 per GPU, achieving an effective batch size of 512. We ensure that this batch size is maintained across our study. We train each model within the DINO framework [3] for 100 epochs. The learning rate initiates at 0 and linearly increases over the initial 10 epochs to reach its baseline value of $lr = 0.0005$. Following this initial warm-up, we decay the learning rate using cosine scheduling [31]. We also use weight decay with an initial value of 0.04, which is increased using a cosine scheduler [31] to 0.4. For the second training phase, we freeze each model as the encoder and train a single fully connected layer using cross-entropy loss. We then use the predictions to measure the accuracy of classification performance on the balanced dataset.

In our experiments, we deploy two versions of Vision Transformer (ViT) encoders: ViT-tiny (ViT-ti) with an embedding dimension of 192 and ViT-Small (ViT-S) with an embedding dimension of 384, both featuring a depth of 12 layers. As highlighted in the Mamba paper [17], a single transformer layer is equivalent to two Mamba block layers. Consequently, for the Vision Mamba tiny (Vim-ti) configuration, we set the depth to 24 layers with an embedding dimension of 192 to maintain comparability. For an enhanced version of Vim-ti, denoted as ViT-ti-plus, we decreased the

Aggregator Model	Zooming Level	Encoder	Pretrained Dataset	Pre-training Method	# Params (Millions)	ACC	F1-Score	AUC
Max-Pooling	20x	ResNet50	ImageNet	Supervised	25	78.95	71.06	81.28
Mean-pooling	20x	ResNet50	ImageNet	Supervised	25	76.69	70.41	80.07
AB-MIL [26]	20x	ResNet50	ImageNet	Supervised	25	90.06	87.40	94.00
DSMIL [29]	20x	ResNet50	ImageNet	Supervised	25	90.17	87.65	94.57
CLAM-SB [32]	20x	ResNet50	ImageNet	Supervised	25	90.31	87.89	94.65
CLAM-MB [32]	20x	ResNet50	ImageNet	Supervised	25	90.14	88.10	94.70
TransMIL [38]	20x	ResNet50	ImageNet	Supervised	25	89.22	85.10	93.51
DTFD-MIL [47]	20x	ResNet50	ImageNet	Supervised	25	90.22	87.62	95.15
MHIM-MIL [40]	20x	ResNet50	ImageNet	Supervised	25	92.48	90.75	96.49
HIPT [4]	20x	ViT-s	TCGA	DINO	22	NA	NA	95.7
iBot-COAD [14]	20x	ViT-B	TCGA-COAD	iBot	86	NA	NA	94.5
CTransPath [45]	20x	CNN	TCGA+PAIP	Swin	NA	92.2	NA	94.2
CLAM-SB [32]	10x	ViT-ti	Cam16	DINO	6	90.69	86.36	87.60
CLAM-SB [32]	10x	Vim-ti	Cam16	DINO	7	93.02	90.32	95.81
CLAM-SB [32]	10x	Vim-ti-plus	Cam16	DINO	13	93.79	91.83	97.39
CLAM-SB [32]	10x	ViT-s	Cam16	DINO	22	94.57	92.47	96.76
CLAM-SB [32]	10x	Vim-s	Cam16	DINO	26	92.24	89.79	98.85

Table 5. Results of Slide Level Classification on Camelyon16 [1] dataset. We provide some methods from the literature as a baseline. However, it is essential to note that our method specifically pre-trains on the training set of Camelyon16. As a result, our improvements should not be compared to the other methods in the field. Our goal in this study is to compare Vision Transformers vs Vision Mamba. Furthermore, we acquired ACC and F1-Score values without searching for the optimal threshold. Hence, AUC is the primary metric to compare the performances of these two methods.

depth to 12 layers while increasing the embedding dimension to 384. This modification aims to investigate the impact of the embedding dimension size on the performance of Vim models. For the Vision Mamba small (Vim-S) setup, we select an embedding dimension of 384 and a depth of 24 layers. The list of trainable parameters across these configurations is provided in Table 4. For slide-level classification, we use the CLAM [32] pipeline to extract patches without using ROI labels. Then, we pass those patches to our encoder model and get the corresponding features. Subsequently, we use the extracted features in the MIL pipelines to classify the whole slides without using the ROI regions.

5.3. Results

In this section, we report the results of our experiments. We first report the performance of both ViT and Vim models for slide-level classification using MIL-based [11] methods. Subsequently, we report the performance on patch-level classification on PCam-224 datasets extracted using the ROI labels of Camelyon16.

Slide-Level Classification For slide-level classification, we use the pre-processing pipeline of CLAM [32] to extract non-background patches from each slide. Then, we use each model as a feature extractor to compute an embedding for that given patch. Subsequently, we use CLAM-SB [32] to classify the slide based on the extracted patch-level feature embeddings. We provide a comparison between our method

and other methods on Camelyon16 in Table 5. For the baseline results, we acquired them from [39], [14], and [45]. We report the F1 Score and Accuracy without searching for an optimal threshold. Hence, while we report them, the ROC AUC is the metric to consider for comparison. The results indicate a substantial improvement in performance with the Vim-ti model, which advances the AUC from 87.60 (as achieved by the ViT-ti model) to 95.81. This increase is particularly significant given the comparable number of parameters involved. This significant enhancement underscores the efficacy of the Vim model in more resource-constrained parameter settings. When considering scaled models, the Vim-ti-plus model, with 13 million parameters, outperforms the 22 million ViT-S model by achieving an AUC of 97.39, compared to the latter’s 96.76. This suggests that the Vim model not only competes effectively at larger scales but can also offer efficient improvements. Furthermore, by scaling up to the Vim-s model with 26 million parameters, we see it remains competitive and sets the highest AUC of 98.85 among all models, emphasizing the Vim architecture’s consistent performance advantages at various scales. This pattern suggests that Vim models excel compared to ViT models in smaller-scale applications and maintain that level of performance when the number of parameters is increased.

Patch-Level Classification While the main goal of the Camelyon16 dataset is slide-level classification, we also provide the performance of models on the extracted patch-

level datasets PCam-224-5x and PCam-224-10x to showcase the learning capability of the proposed method on weakly-supervised classification tasks. For evaluating the patch-level classification performance, we use accuracy for both PCam-224 datasets since they are balanced. The results for PCam-224-10x are provided in Table 6. We provide random weights not as a comparison metric but to show the effect of self-supervised learning in the given task. Our observations indicate that the Vim-ti model achieves superior performance over the ViT-ti of a comparable number of parameters. In contrast, ViT-s outperforms the Vim-ti-plus and Vim-S models by a small margin in this task. Furthermore, the performance improvement for the given task from Vim-ti-plus to Vim-S is minimal, while the number of parameters doubles.

Weights	Model	Linear Evaluation (ACC)
Random Weights	ViT-ti	79.16
	Vim-ti	74.87
	Vim-ti-plus	77.29
	ViT-s	83.19
Cam16-10x DINO	ViT-ti	94.59
	Vim-ti	95.99
	Vim-ti-plus	96.48
	ViT-s	96.65
	Vim-s	96.51

Table 6. Results on PCam-224-10x.

We also perform some tests for 5x zooming level using PCam-224-5x. The results are summarized in the Table 7. We provide random weights not as a comparison metric but to show the effect of self-supervised learning in the given task. We have a consistent result as before, with Vim-ti outperforming its alternative ViT-ti with a similar number of parameters. Furthermore, Vim-ti-plus also outperforms ViT-s with approximately half the parameters.

Weights	Model	Linear Evaluation (ACC)
Random Weights	ViT-ti	82.08
	Vim-ti	77.12
	Vim-ti-plus	78.37
	ViT-s	81.61
Cam16-5x DINO	ViT-ti	88.87
	Vim-ti	90.45
	Vim-ti-plus	90.39
	ViT-s	90.36

Table 7. Results on PCam-224-5x.

5.4. Explainability

We use GradCam [37] to generate heatmaps of areas that each model is attending to the most by considering the activation maps of the CLS token of each model with respect to the input. Note that while we use the CLS token, both models are trained without labels.

Our board-certified pathologist reviewed these CLS token activation heat maps generated from each model on a representative subset of tumor patches and compared Vim-based models against ViT-based models. Vim heatmaps were generally oriented foremost to distinctive cancer-specific cellular features and the interface with non-cancer cells. Concurrently, ViT heatmaps foremost highlighted atypical cancer cells. For example, Vim heatmaps highlighted intracellular mucin, a feature nearly 100% specific for cancer, as well as adjacent activated lymphocytes, which are biologically reactive to the presence of cancer in the lymph node (Fig 4). Indeed, both of these features are exclusively found in the presence of cancer.




Figure 4. Representative tumor patch with Vim-s heatmap. The red asterisks highlight intracellular mucin in cancer cells. The yellow asterisks highlight stromal features adjacent to cancer cells. (The heatmaps are generated at 10x and overlaid on 40x images.)




Figure 5. Representative tumor patch with ViT-s heatmap. The red asterisks highlight areas centralized on cancer cells. The yellow asterisks highlight other features, notably a focus of intracellular mucin (top-right) and a stromal cell (middle). (The heatmaps are generated at 10x and overlaid on 40x images.)

Alternatively, ViT heatmaps were focused on cancer cells as a whole, notably very atypical cancer cells, nearly in a dichotomous fashion of cancer versus non-cancer (Fig 5).

Overall, while both models worked efficiently, Vim could identify biological features that pathologists use to validate the presence of cancer inside lymph nodes, while ViT was oriented toward distinguishing whole cancer cells from non-cancer cells. This is an indication that the proposed Vim4Path framework mimics the workflow of pathologists.

5.5. Ablation Study

In this section, we provide an ablation study on the impact of zooming level for both slide-level and patch-level classification. Table 8 shows the impact of transferring a weight trained on one zooming level to another. We notice that doing MIL at the 5x zooming level performs poorly, while MIL at a 10x zooming level gives reasonable results. Furthermore, a model pretrained on 5x does not transfer well to 10x for feature extraction. This performance drop is expected as the data distribution shifts and the fact that the models trained at 5x had fewer pre-training datasets, as summarized in Table 2. The only model that performs reasonably with distribution shift is ViT-s trained on 10x when tested on 5x.

Pretrained Weights	Model	MIL for Cam16 at 5x		MIL for Cam16 at 10x	
		AUC	ACC	AUC	ACC
Cam16-5x DINO	ViT-ti	65.91	73.64	68.85	65.11
	Vim-ti	71.30	73.64	64.41	69.76
	Vim-ti*	79.03	75.96	74.08	72.86
	ViT-s	67.72	76.74	69.87	74.41
Cam16 DINO-10x	ViT-ti	75.05	72.86	87.60	90.69
	Vim-ti	62.14	54.26	95.81	93.02
	Vim-ti*	71.47	79.06	97.39	93.79
	ViT-s	84.89	81.39	96.76	94.57
	Vim-s	79.64	79.06	98.85	92.24

Table 8. Effect of Zooming level on MIL. * stands for Vim-ti-plus

Table 9 displays the performance of patch-level classification across different zooming levels. It reveals that feature representations learned at one magnification can be applied to other levels in patch classification tasks. Most importantly, models pretrained with a larger dataset at a 10x zooming level (summarized in Table 2) demonstrate superior representation quality during linear evaluation at a 5x level compared to those trained at 5x. This discrepancy shows the significance of dataset size in the pre-training phase in enhancing the quality of the learned representations.

5.6. Discussion

The Vim model’s ability to outperform the ViT at smaller scales is significant. Vim’s architecture is potentially more

Pretrained Weights	Model	Linear Eval (5x)	Linear Eval (10x)
Cam16-5x DINO	ViT-ti	88.87	89.29
	Vim-ti	90.45	91.03
	Vim-ti*	90.39	88.72
	ViT-s	90.36	89.92
Cam16-10x DINO	ViT-ti	90.99	94.59
	Vim-ti	93.01	95.99
	Vim-ti*	93.23	96.48
	ViT-s	93.00	96.65
	Vim-s	93.04	96.51

Table 9. Effect of Zooming level on patch classification. * stands for Vim-ti-plus

efficient, harnessing fewer parameters to achieve higher accuracy. This efficiency is crucial in scenarios with limited computational resources where cost considerations are essential. Vim is also comparable to ViT on larger scales. Furthermore, the Vim seems to emulate the pathologists’ workflow. These findings suggest that Vim is a promising architecture design for pathology applications that aligns more with clinical workflows. However, our study has limitations. The compute constraints prevented exhaustive hyperparameter tuning and more extended training regimes, which could unlock even higher performance from the Vim model. Additionally, the performance improvements at different zooming levels raise essential questions about the transferability of features across scales and the role of dataset size in pre-training. Future work should extend the comparative analysis to include more diverse datasets. There is also a need for deeper investigation into the interpretability of the models to build trust in deploying these models in clinical settings.

6. Conclusion

In this paper, we studied the utility of the Vision Mamba (Vim) encoder model within a Self-Supervised Learning (SSL) framework, particularly DINO, to extract feature embeddings for computational pathology. In particular, we have benchmarked the Camelyon16 dataset to extract image patches from WSIs without labels for training the Vim encoder using the DINO framework. We modified the Vim model to intake arbitrary input image size using positional embedding interpolation, making the Vim model adaptable within DINO for SSL. We compare Vim with ViT in patch-level and slide-level classification tasks on the Camelyon16 dataset, demonstrating Vim’s superior performance at small scales. We also show that Vim remains competitive and often outperforms ViT at larger scales. Additionally, our findings reveal that Vim, in contrast to ViT, replicates the process pathologists use to identify cancer. Our study indicates the potential benefits of employing state space models in computational pathology.

References

- [1] 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. [2](#), [5](#), [6](#)
- [2] Mikhail Binnewies, Edward W Roberts, Kelly Kersten, Vincent Chan, Douglas F Fearon, Miriam Merad, Lisa M Coussens, Dmitry I Gabrilovich, Suzanne Ostrand-Rosenberg, Catherine C Hedrick, et al. Understanding the tumor immune microenvironment (time) for effective therapy. *Nature medicine*, 24(5):541–550, 2018. [4](#)
- [3] 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. [1](#), [2](#), [3](#), [5](#)
- [4] 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, 2022. [2](#), [6](#)
- [5] 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. [2](#)
- [6] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E Hinton. Big self-supervised models are strong semi-supervised learners. *Advances in neural information processing systems*, 33:22243–22255, 2020. [2](#)
- [7] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. [2](#)
- [8] 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, 2021. [2](#)
- [9] Ozan Ciga, Tony Xu, and Anne Louise Martel. Self supervised contrastive learning for digital histopathology. *Machine Learning with Applications*, 7:100198, 2022. [2](#)
- [10] 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, 2009. [1](#), [2](#)
- [11] Thomas G. Dietterich, Richard H. Lathrop, and Tomás Lozano-Pérez. Solving the multiple instance problem with axis-parallel rectangles. *Artificial Intelligence*, 89(1):31–71, 1997. [1](#), [2](#), [6](#)
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. [2](#), [3](#)
- [13] Zijie Fang, Yifeng Wang, Zhi Wang, Jian Zhang, Xiangyang Ji, and Yongbing Zhang. Mammil: Multiple instance learning for whole slide images with state space models. *arXiv preprint arXiv:2403.05160*, 2024. [2](#)
- [14] Alexandre Filiot, Ridouane Ghermi, Antoine Olivier, Paul Jacob, Lucas Fidon, Alice Mac Kain, Charlie Saillard, and Jean-Baptiste Schiratti. Scaling self-supervised learning for histopathology with masked image modeling. *medRxiv*, pages 2023–07, 2023. [6](#)
- [15] 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. [2](#)
- [16] Daniel Y Fu, Tri Dao, Khaled K Saab, Armin W Thomas, Atri Rudra, and Christopher Ré. Hungry hungry hippos: Towards language modeling with state space models. *arXiv preprint arXiv:2212.14052*, 2022. [2](#)
- [17] Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*, 2023. [2](#), [3](#), [4](#), [5](#)
- [18] Albert Gu, Tri Dao, Stefano Ermon, Atri Rudra, and Christopher Ré. Hippo: Recurrent memory with optimal polynomial projections. *Advances in neural information processing systems*, 33:1474–1487, 2020. [2](#)
- [19] Albert Gu, Karan Goel, and Christopher Ré. Efficiently modeling long sequences with structured state spaces. *arXiv preprint arXiv:2111.00396*, 2021. [2](#)
- [20] Albert Gu, Karan Goel, and Christopher Ré. Efficiently modeling long sequences with structured state spaces. *arXiv preprint arXiv:2111.00396*, 2021. [2](#)
- [21] Albert Gu, Karan Goel, Ankit Gupta, and Christopher Ré. On the parameterization and initialization of diagonal state space models. *Advances in Neural Information Processing Systems*, 35:35971–35983, 2022. [2](#)
- [22] 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. [1](#)
- [23] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. [2](#)
- [24] 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. [2](#)
- [25] Mahdi S Hosseini, Babak Ehteshami Bejnordi, Vincent Quoc-Huy Trinh, Lyndon Chan, Danial Hasan, Xingwen Li, Stephen Yang, Taehyo Kim, Haochen Zhang, Theodore Wu, et al. Computational pathology: a survey review and the way forward. *Journal of Pathology Informatics*, page 100357, 2024. [1](#), [2](#)

- [26] Maximilian Ilse, Jakub Tomczak, and Max Welling. Attention-based deep multiple instance learning. In *International conference on machine learning*, pages 2127–2136. PMLR, 2018. 2, 6
- [27] Maximilian Ilse, Jakub Tomczak, and Max Welling. Attention-based deep multiple instance learning. In *International conference on machine learning*, pages 2127–2136. PMLR, 2018. 2
- [28] 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. 1
- [29] Bin Li, Yin Li, and Kevin W Eliceiri. Dual-stream multiple instance learning network for whole slide image classification with self-supervised contrastive learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14318–14328, 2021. 2, 6
- [30] Bin Li, Yin Li, and Kevin W Eliceiri. Dual-stream multiple instance learning network for whole slide image classification with self-supervised contrastive learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14318–14328, 2021. 2
- [31] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. 5
- [32] Ming Y Lu, Drew FK Williamson, Tiffany Y Chen, Richard J Chen, Matteo Barbieri, and Faisal Mahmood. Data-efficient and weakly supervised computational pathology on whole-slide images. *Nature Biomedical Engineering*, 5(6):555–570, 2021. 2, 3, 5, 6
- [33] Jesper Molin, Morten Fjeld, Claudia Mello-Thoms, and Claes Lundström. Slide navigation patterns among pathologists with long experience of digital review. *Histopathology*, 67(2): 185–192, 2015. 4
- [34] Lambda Moses and Lior Pachter. Museum of spatial transcriptomics. *Nature methods*, 19(5):534–546, 2022. 4
- [35] Michael Poli, Stefano Massaroli, Eric Nguyen, Daniel Y Fu, Tri Dao, Stephen Baccus, Yoshua Bengio, Stefano Ermon, and Christopher Ré. Hyena hierarchy: Towards larger convolutional language models. *arXiv preprint arXiv:2302.10866*, 2023. 2
- [36] Daniela F Quail and Johanna A Joyce. Microenvironmental regulation of tumor progression and metastasis. *Nature medicine*, 19(11):1423–1437, 2013. 4
- [37] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017. 7
- [38] 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. 2, 6
- [39] Wenhao Tang, Sheng Huang, Xiaoxian Zhang, Fengtao Zhou, Yi Zhang, and Bo Liu. Multiple instance learning framework with masked hard instance mining for whole slide image classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4078–4087, 2023. 6
- [40] Wenhao Tang, Sheng Huang, Xiaoxian Zhang, Fengtao Zhou, Yi Zhang, and Bo Liu. Multiple instance learning framework with masked hard instance mining for whole slide image classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4078–4087, 2023. 6
- [41] Manuel Tran, Sophia J Wagner, Melanie Boxberg, and Tingying Peng. S5cl: unifying fully-supervised, self-supervised, and semi-supervised learning through hierarchical contrastive learning. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 99–108. Springer, 2022. 2
- [42] Aäron Van Den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *International conference on machine learning*, pages 1747–1756. PMLR, 2016. 4
- [43] Bastiaan S Veeling, Jasper Linmans, Jim Winkens, Taco Cohen, and Max Welling. Rotation equivariant cnns for digital pathology. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2018: 21st International Conference, Granada, Spain, September 16–20, 2018, Proceedings, Part II 11*, pages 210–218. Springer, 2018. 5
- [44] Alex D Waldman, Jill M Fritz, and Michael J Lenardo. A guide to cancer immunotherapy: from t cell basic science to clinical practice. *Nature Reviews Immunology*, 20(11): 651–668, 2020. 4
- [45] 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. 6
- [46] Shu Yang, Yihui Wang, and Hao Chen. Mambamil: Enhancing long sequence modeling with sequence reordering in computational pathology. *arXiv preprint arXiv:2403.06800*, 2024. 2
- [47] Hongrun Zhang, Yanda Meng, Yitian Zhao, Yihong Qiao, Xiaoyun Yang, Sarah E Coupland, and Yalin Zheng. Dtf-d-mil: Double-tier feature distillation multiple instance learning for histopathology whole slide image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18802–18812, 2022. 2, 6
- [48] Lianghui Zhu, Bencheng Liao, Qian Zhang, Xinlong Wang, Wenyu Liu, and Xinggang Wang. Vision mamba: Efficient visual representation learning with bidirectional state space model. *arXiv preprint arXiv:2401.09417*, 2024. 1, 2, 3