

Improving Representation Learning for Histopathologic Images with Cluster Constraints

Weiye Wu¹ Chongyang Gao² Joseph DiPalma¹ Soroush Vosoughi¹ Saeed Hassanpour¹

¹Dartmouth College ²Northwestern University

{weiyi.wu.gr,joseph.r.dipalma.gr,soroush.vosoughi,saeed.hassanpour}@dartmouth.edu
chongyanggao2026@u.northwestern.edu

Abstract

Recent advances in whole-slide image (WSI) scanners and computational capabilities have significantly propelled the application of artificial intelligence in histopathology slide analysis. While these strides are promising, current supervised learning approaches for WSI analysis come with the challenge of exhaustively labeling high-resolution slides—a process that is both labor-intensive and time-consuming. In contrast, self-supervised learning (SSL) pretraining strategies are emerging as a viable alternative, given that they don't rely on explicit data annotations. These SSL strategies are quickly bridging the performance disparity with their supervised counterparts. In this context, we introduce an SSL framework. This framework aims for transferable representation learning and semantically meaningful clustering by synergizing invariance loss and clustering loss in WSI analysis. Notably, our approach outperforms common SSL methods in downstream classification and clustering tasks, as evidenced by tests on the Camelyon16 and a pancreatic cancer dataset. The code and additional details are accessible at <https://github.com/wwyi1828/CluSiam>.

1. Introduction

Histopathology slide analysis remains the gold standard for cancer diagnosis and prognosis. In recent years, researchers have seen the burgeoning adoption of digital histopathology slides in pathology laboratories, thanks to the availability of digital pathology scanners and advancements in computer vision, revolutionizing computational pathology [24]. While adoption of digital slides has accelerated, progress has been hindered by the exceptionally high resolution of whole slide images (WSIs), often exceeding $40,000 \times 40,000$ pixels, which makes directly applying

standard computer vision models to WSIs not feasible. Furthermore, downsampling WSIs to a more manageable magnification level results in a substantial loss of fine-grained visual information.




Figure 1: The framework of CluSiam. View 1 and View 2 are distinct augmentations of the same images, pooled together for clustering. The invariance loss (solid line) aligns representations of the two views, while the cluster loss (dashed line) pushes cluster centroids apart.

To address the challenges, WSIs are commonly subdivided into more manageable patches through sliding-window techniques. These patches are then labeled using annotations, forming training data for a patch-level classifier. Features extracted by the trained patch-level classifiers are aggregated to infer slide-level label [21, 42, 43, 40, 33]. However, this annotation-reliant approach has a significant drawback. It's heavily dependent on precise annotations, which are expensive to obtain. Annotating WSIs is a painstaking and error-prone task that requires pixel-by-pixel scrutiny from highly skilled pathologists. The elusive

borders between different tissue patterns introduce variability among pathologists. Additionally, tissue morphology’s inherent variability often further diminishes the accuracy of annotations. Therefore, obtaining precise and consistent annotations remains an uphill battle, even with substantial expertise and time invested by trained pathologists. Inaccurate annotations can potentially lead to inaccurate and inconsistent WSI analysis models [39, 1].

To mitigate the impact of inaccurate annotations, noise-aware learning models have emerged. These methods improve the performance of patch-level feature extractors by either filtering or down-weighting the noisy patches [1, 28, 13]. However, these models are still constrained by the annotation bottleneck. Even acquiring noisy annotations for WSIs demands substantial time and expertise, which motivates the need for annotation-free techniques. They reduce costs and save time. They also eliminate the impact of inaccurate annotations.

In this challenging landscape, annotation-free techniques have emerged as a promising solution. By requiring only whole-slide labels, they not only cut costs and save time but also eliminate the effects of annotation inaccuracies. Among these, Chen *et al.* [9] proposed a method that leverages a unified memory mechanism to train convolutional neural networks (CNNs) directly with numerous images. However, this approach is constrained to lower magnification levels, restricting pixel sizes to above 2 μm . Conversely, other studies [33, 30, 14, 15] have shown that achieving better results is possible by employing higher or multi-scale magnification levels across a variety of model designs.

Weakly supervised techniques have gained popularity as an annotation-free approach that retains high-resolution details by utilizing slide-level labels instead of exhaustive patch-level annotations. Obtaining slide-level labels is less laborious compared to exhaustive patch-level annotations. Thus, weakly supervised learning has become particularly popular for histology slide classification tasks [36, 37, 46]. These methods employ slide-level labels as weak supervision for all patches within a slide. Multiple instance learning (MIL) models leverage this by treating slides as positive or negative bags, with patches as instances [29, 30, 22, 46, 3, 34]. However, MIL models have some limitations. They often neglect important contextual cues across a whole slide. Additionally, off-the-shelf feature extractors pretrained on natural images fail to sufficiently capture tissue morphology. These drawbacks motivate exploring self-supervised approaches for histology slides.

Self-supervised learning (SSL) enables models to learn feature representations without the need for labels. SSL methods are rapidly closing the performance gap with supervised approaches. However, SSL typically requires a large sample size. but this is mitigated for high-resolution

histopathology images by splitting WSIs into numerous small patches. In computational pathology, self-supervision methods become an appealing solution for annotation-free WSI analysis [30, 27, 26, 44, 7]. These methods utilize multiple-instance learning to aggregate self-supervised patch representations. They have demonstrated the capability to match the performance of state-of-the-art supervised methods while reducing the annotation burden on pathologists by eliminating the need for manual annotations.

One of the key paradigms of SSL is contrastive-based SSL [19, 11, 35, 38, 25]. They may not be the most effective in histopathology image analysis because adjacent patches from a WSI can be very similar in their morphological features, making them unsuitable as negative sample pairs. These methods also rely on a large number of negative pairs. To avoid the need for negative pairs, some knowledge-distillation-based methods [12, 6, 18] concentrate solely on positive sample pairs, which are defined using augmented views of the same image. However, only focusing on positive pairs might prevent them from learning global information, as their objective functions only consider augmentations from the same image.

Apart from SSL representation learning, another pivotal technique gaining attention is clustering. Clustering is an unsupervised learning approach where similar samples are grouped to ensure intra-cluster cohesion and inter-cluster separation. In the domain of WSI retrieval, clustering could be instrumental. Wang *et al.* [41] employed a K-Means clustering-driven architecture, while Chen *et al.* [8] integrated a self-supervised variational autoencoder with the K-Means algorithm, both for WSI retrieval systems. Given the growing prominence of such methods in WSI retrieval, there’s an increasing demand to refine these clustering algorithms within computational pathology. Clustering shares similarities with representation learning. This has inspired clustering-based SSL methods that use pseudo-labels from iterative K-Means clustering algorithms for training feature encoders. Although these methods can learn effective image representations, they may not improve the performance of the actual clustering tasks as they cluster images into thousands of groups, which might hinder their direct use for histopathology image retrieval. Large cluster counts make identifying relevant groups challenging.

To address the shortcomings, we propose Cluster-Siam (CluSiam), a framework that decouples clustering from representation learning and retains only the most relevant and interpretable clusters for medical applications. CluSiam takes advantage of an existing self-supervised backbone to extract representations. We introduce a cluster loss to guide the backbone in learning effective representations while generating accurate, interpretable cluster assignments for histopathology images. (Figure 1). Our experiments demonstrate that CluSiam outperforms baselines on down-

stream classification tasks. Additionally, our adaptive clustering algorithm outperforms K-Means in clustering, resulting in improved cluster assignments. In addition, our cluster assigner emerges as a by-product of the representation learning process, thus introducing only a small additional computational cost once the training is complete.

The contributions of this paper can be summarized as follows:

- We propose CluSiam, a SSL framework for image representation learning and clustering that combines invariance loss and cluster loss (Figure 2)
- We compare the performance of different SSL frameworks and demonstrate that CluSiam outperforms popular SSL methods on multiple histopathology datasets.
- CluSiam provides an efficient and accurate way to cluster histopathology images without either patch-level annotations or slide-level labels, with clustering performance better than the widely used K-Means clustering in digital pathology.




Figure 2: The details of the CluSiam framework. The invariance loss maximizes the on-diagonal elements of the similarity matrix between views. The cluster assigner takes the concatenated views as input and generates cluster assignments. The cluster loss minimizes the off-diagonal elements of the similarity matrix between cluster centroids.

2. Related Works

2.1. Self-supervised learning

Self-supervised learning (SSL) methods have recently demonstrated effectiveness for computer vision tasks by learning representations without reliance on manual labels. SSL techniques leverage the intrinsic structure and consistency of the data itself as a supervisory signal. Several paradigms have arisen, including contrastive-based,

knowledge-distillation-based, clustering-based, and information maximization-based approaches. Typically, these techniques function by generating augmented pairs of views from a single data instance and directing the model to produce similar outputs for each view.

Contrastive learning stands out as a key self-supervised approach in representation learning, with the goal of deriving informative and concise representations from unstructured data. A slew of methods grounded in contrastive learning have been proposed, including Contrastive Predictive Coding (CPC) [35], SimCLR [11], MoCo [19], and NNCLR [17]. CPC, recognized for its widespread application, employs an autoregressive model to predict future observations based on past observations, rendering it particularly apt for sequential data. MoCo and SimCLR, two other popular contrastive learning techniques for instance discrimination, generate positive pairs by utilizing two different views (augmentations) of the same image and negative pairs by pairing augmentations of different data points. The principle behind contrastive learning is to distance negative pairs while converging positive pairs. However, achieving optimal performance with these methods often necessitates a plethora of negative pairs. MoCo addresses this issue by implementing momentum encoders and a memory bank mechanism, while SimCLR capitalizes on large batch sizes for negative pair comparisons. In NNCLR, the model learns representations by minimizing the distance between an anchor and its nearest neighbor in the momentum encoder’s output space while maximizing the distance to other negative samples. This method streamlines the utilization of negative samples in a batch, curtailing the need for large batch sizes and memory banks while maintaining competitive performance compared to other contrastive learning methods.

Knowledge-distillation-based methodologies, such as BYOL [18] and SimSiam [12], aim to enhance performance with smaller batch sizes and without the need for negative samples. In stark contrast to their contrastive counterparts, these non-contrastive techniques employ only positive pairs, eliminating the demand for large batch sizes or a memory bank mechanism. BYOL stands out with its momentum update mechanism, which renders negative pairs unnecessary. It establishes a target network by applying an exponential moving average to the online network’s weights. This “moving target” offers the online network a stable benchmark to aim for during training, pushing the network away from trivial solutions and encouraging richer representations. SimSiam, influenced by BYOL, streamlines the process by forgoing the moving target. Instead, it adopts a symmetric architecture, where dual networks reciprocally predict each other’s outputs. To ensure the representations are non-trivial, SimSiam utilizes a stop-gradient operation. Nonetheless, despite their batch efficiency, both

BYOL and SimSiam remain susceptible to collapsing into trivial solutions.

Clustering-based SSL methods, such as DeepCluster [4] and Prototypical Contrastive Learning (PCL) [31], offer an innovative angle to representation learning by capitalizing on iterative pseudo-labeling to cluster the learned representations. DeepCluster updates its network parameters according to pseudo-labels produced by the K-Means clustering algorithm. This aligns the network more closely with the inherent data distribution. Expanding on the foundation laid by MoCo, PCL integrates the ProtoNCE loss and K-Means clustering, aiming to refine image embeddings by nudging them closer to their respective prototypes by optimizing the ProtoNCE loss function. SwAV [5], another clustering-based SSL method, shares similarities with SimSiam architecture but differentiates itself with a swap prediction mechanism. Specifically, SwAV aligns the cluster assignments of one augmentation with the representations of another version of the same image. These assignments are fine-tuned using the Sinkhorn algorithm, ensuring a balanced distribution. In embracing this strategy, SwAV learns invariant features that encapsulate crucial semantic information.

Information maximization-based SSL approaches strive to learn representations by maximizing the invariance of corresponding features while minimizing the covariance between different features. Noteworthy methods in this realm include VICReg [2] and Barlow Twins [45]. VICReg specifically enforces feature representation invariance by amplifying their variance and curtailing the covariance between different features. This strategy ensures that the learned representations are not only informative but also capture the core attributes of the data. Contrasting with VICReg, which explicitly maximizes the variance of individual features, Barlow Twins centers its focus on minimizing cross-correlations across feature dimensions. It achieves this by mitigating the cross-correlation between the outputs of twin networks, each processing a distinct augmentation of the same image while concurrently accentuating the invariance of matching features.

2.2. SSL in Pathology

The advent of whole-slide scanners has enabled the digitization of histopathological slides, gradually transforming the field of anatomical pathology into a data-abundant domain. Recognizing this, SSL techniques are being increasingly employed in computational pathology to take advantage of the abundance of unlabeled data. In conjunction with smaller labeled datasets, these techniques promise to elevate diagnostic precision and bolster predictive modeling.

In the realm of computational pathology, SSL is gradually gaining adoption as a means to tackle challenges such

as acquiring annotations for pathology slides, managing high-resolution images, and addressing the substantial variability in their features. Numerous studies have leveraged SSL for extracting features from WSIs and have utilized these features to achieve promising results in downstream histopathology image analysis tasks [10, 16, 26, 30]. With the escalating adoption of SSL in computational pathology, there is a growing necessity to determine the applicability of general self-supervised methods to histopathology images. A recent benchmarking study [27] gauged various SSL methods across diverse pathology datasets for various downstream tasks, such as classification and nuclei instance segmentation tasks. Their results indicate that SSL can considerably uplift the performance on downstream tasks on histopathology images compared to ImageNet pre-trained and supervised models, especially when labeled data is scarce.

3. Method

We recap SimSiam and then present our method for self-supervised representation learning and clustering.

3.1. Preliminaries: SimSiam

Self-supervised visual representation learning is a method for learning an embedding function that maps an input image x to a representation. This is typically achieved by using a similarity measure designed to enforce similarity between augmented views. Starting with sets of data transformations, \mathcal{T}_1 and \mathcal{T}_2 , we randomly sample transformations $t_1, t_2 \sim \mathcal{T}_1, \mathcal{T}_2$ and produce augmented views $x_1 = t_1(x)$ and $x_2 = t_2(x)$. An encoder f is used to produce representations $y_1 = f(x_1)$ and $y_2 = f(x_2)$, which are then fed to a projector h to produce projections $z_1 = h(y_1)$ and $z_2 = h(y_2)$. As in SimSiam, we pass z_1 through a predictor g to produce the prediction $p_1 = g(z_1)$. Additionally, we swap the views and produce a symmetric loss as follows:

$$\mathcal{L}_{inv} = \text{sim}(p_1, \text{sg}(z_2)) + \text{sim}(p_2, \text{sg}(z_1)) \quad (1)$$

where

$$\text{sim}(p, z) = -\frac{p}{\|p\|_2} \cdot \frac{z}{\|z\|_2}, \quad (2)$$

$\|\cdot\|$ is the ℓ_2 -norm, and $\text{sg}(\cdot)$ is the stop-gradient operation to prevent collapse.

3.2. CluSiam: Cluster-Constrained SSL

We build upon the SimSiam architecture by adding a cluster assigner q that operates on the projections produced by h . We use the outputs of h as the inputs to q because the batch normalization layers (\mathcal{BN}) in h stabilize the distribution of q 's inputs. We also introduce \mathcal{BN} in q because controlling the input scale is crucial for generating cluster

assignment probabilities using softmax. Given batches of views X_1, X_2 , we produce projections $Z_1 = \{h(f(x_i^{(1)})) : x_i^{(1)} \in X_1\}, Z_2 = \{h(f(x_i^{(2)})) : x_i^{(2)} \in X_2\}$. We then concatenate these projections to obtain $Z = \text{concat}(Z_1, Z_2)$. The cluster assigner q maps the concatenated projections Z to cluster representations A defined as:

$$A = \frac{\exp(A_i)/\tau}{\sum_i \exp(A_i)/\tau} \in \mathbb{R}^{2 \cdot N \times K} \quad (3)$$

where A_{ij} is the element at the i -th row and j -th column of A , τ is the temperature, and K is the exploration space that represents the maximum number of clusters allowed during the training. This operation is applied along each row, meaning that for every data point i , the sum of A_{ij} over all clusters, j 's, equals 1. Finally, we map the cluster representation to the clusters as follows:

$$C = \text{nonzero} \left(\frac{\text{argmax}(A^T) \cdot \text{sg}(Z)}{\|\text{argmax}(A^T) \cdot \text{sg}(Z)\|_2} \right) \in \mathbb{R}^{k \times D}. \quad (4)$$

Notably, $\text{argmax}(A)$ is not differentiable, so there is no real gradient for this operation. We approximate the gradient similar to the straight-through estimator and just copy gradients from A to $\text{argmax}(A)$, $\nabla_{\text{argmax}(A)} C = \nabla_A C$, making it possible for backpropagation. $\text{sg}(\cdot)$ denotes the stop-gradient operation, $\text{nonzero}(\cdot)$ filters out vectors along the row dimension that are all zeros, and k represents the count of non-zero centroids. The dimension D corresponds to the feature dimension and is consistent with the dimensions of p and z .

In our clustering module, unlike common clustering methods that use inter-class similarity or other SSL techniques that focus on the invariance between two different views, we do not impose any restrictions on inter-class similarity or the assignments between two different views, a_i and a_{n+i} . Our cluster loss is solely defined by inter-cluster separation. This separation is defined as:

$$\mathcal{L}_{cluster} = - \frac{\sum_i \sum_{i \neq j} (C^T C)_{ij}}{2C_k^2}. \quad (5)$$

Importantly, all the tensors in (5) are ℓ_2 normalized, so equation (5) can be interpreted as representing the average cosine similarities between clusters. Furthermore, the stop-gradient operation is applied to all elements in (4) to prevent the cluster assigner from collapsing to a trivial solution where all samples are assigned to the same set or cluster. We can view $A = a_1, \dots, a_{2n}$ as a latent variable, with our goal being to minimize $\mathcal{L}(A, Z)$. This optimization problem can be solved by an alternating algorithm that fixes one set of variables and solves for the other set.

The intra-class similarity term was not introduced in this design because we neither use a contrastive formula-

tion loss function like SimCLR nor an additional projection head to introduce the knowledge distillation architecture like BYOL. Directly optimizing for high intra-class similarity is prone to collapsing all samples into the same representation, which is a trivial solution and does not capture any meaningful information.

However, our cluster learning is still prone to collapse due to the presence of the softmax and argmax functions. The hard cluster assignment of the softmax and argmax functions limits the model to updating only the maximum scoring cluster during training. This does not encourage exploration of different combinations of cluster assignments, and other potential cluster assignments are left out of the backpropagation updates. As a result, the model is prone to getting stuck in a trivial single-cluster solution during training, a phenomenon known as ‘‘collapsing’’. This is especially likely as the number of valid clusters decreases. Using softmax and argmax based cluster assignment can exacerbate this problem, leading to the rapid collapse of a single cluster. This prevents the model from learning the intended cluster structures and, in turn, leads to meaningless representations since the cluster loss cannot be effectively optimized with only a single active cluster.

The task of cluster assignment can be viewed as a decision-making process in which the cluster assigner determines which samples belong in the same cluster. Striking a balance between exploration and exploitation is crucial in this decision-making process. Exploration refers to trialing different actions to learn more about the environment and their associated losses, while exploitation refers to choosing the action currently known to possess the lowest expected loss. To prevent the collapse of clustering due to continual updates only to the neuron with the highest probability value, we introduce randomness into the decision-making process by adding Gumbel noise [23]. Gumbel noise is a random variable sampled from a Gumbel distribution. It proves useful in discrete action spaces, where a model must choose between a finite set of actions, as in our cluster assignment task. By replacing (3) with (6), the cluster assigner can explore different cluster combinations based on their probabilities, thereby learning more about different cluster combinations and their associated losses.

$$A = \frac{\exp(a_i + g_i)/\tau}{\sum_{i=1} \exp(a_i + g_i)/\tau}, \quad g_i \sim G(0, 1), \quad (6)$$

where G represents the Gumbel distribution.

Therefore, CluSiam can be trained effectively using the composite loss function, as shown in equation (7). In this equation, α serves as a hyperparameter, adjusting the magnitude of the weights. In our implementation, we set $\alpha = 0.5$ by default.

$$\mathcal{L}_{CluSiam} = (1 - \alpha) \cdot \mathcal{L}_{inv} + \alpha \cdot \mathcal{L}_{cluster} \quad (7)$$

4. Experiments

In our experiments, we evaluated the performance of our proposed model on two clinically relevant whole-slide image datasets: the Camelyon16 [32] and the Pancreatic Cancer dataset [43]. To extract representative image patches from the WSIs, we first removed the background and employed a sliding window technique to generate patches of size 224×224 at a $20 \times$ magnification level ($0.5 \mu\text{m}/\text{pixel}$) from the tissue regions of a slide, with no overlap between patches. The Camelyon16 dataset is a publicly available dataset designed for the task of metastasis detection in breast cancer. It includes two classes, positive and negative slides, and consists of 271 training images and 129 testing images. After applying our patch extraction algorithm, we obtained approximately 2.6 million training and 1.2 million testing patches at $20 \times$ magnification for this dataset. The Pancreatic Cancer dataset includes three classes: negative (background class), neoplastic, and positive. This dataset includes 104 training slides and 39 testing slides, yielding approximately 300,000 training and 83,000 testing patches.

In our study, we compared our proposed CluSiam method against a supervised model and several commonly used SSL architectures as baselines. All SSL models, as well as the supervised model, were trained using a ResNet18 [20] backbone for 50 epochs with a batch size of 512. The hyperparameters for training were set to be as identical as possible to their default values specified in the original studies with comparable settings. The detailed hyperparameter settings can be found in the appendix. It’s important to note that different architectures often incorporate varying image augmentation techniques, optimizers, and hyperparameters. As a result, comparisons between baseline methods may not be entirely fair due to inherent configuration differences. Our approach (CluSiam) used identical hyperparameter settings as SimSiam, which allows for a direct and equitable comparison between these two architectures. As a metric for evaluation, we employed two downstream tasks: clustering and classification.

4.1. Clustering

In the clustering task, we evaluated the performance of clustering algorithms using the Rand Index (RI). The performance was compared using different representations and clustering algorithms (Table 1). Importantly, our cluster assigner’s output is different from traditional methods like K-Means. Unlike K-Means, which provides a hard cluster assignment for each patch, our assigner outputs a K-dimensional vector. This structure allows for more flexibility in generating cluster assignments beyond the typical use of argmax. Such probabilistic outputs grant our method greater adaptability in clustering, especially when contrasted with the rigid assignments derived from K-Means. We visualized the cluster assignments of two WSIs gener-

ated simply using argmax, alongside their respective ground truths, in Figure 3.

Encoder	Cluster	Camelyon16	Pancreatic
SimSiam	K-Means	0.509	0.329
CluSiam	K-Means	0.538	0.357
CluSiam	Assigner	0.897	0.569

Table 1: Rand Index on the testing set.

4.2. Classification

For the classification task, we evaluated model performance using accuracy and Area Under the ROC Curve (AUC) metrics. We aggregated patch-level predictions to slide-level predictions using two multiple-instance learning techniques: Max-Pooling (Max) and Dual-Stream Multiple-Instance Learning (DSMIL). Given that optimal hyperparameters for the MIL models may differ based on the representations learned by various SSL methods, we conducted grid searches over learning rates [$1e-2$, $1e-3$, $1e-4$] and weight decays [$1e-2$, $1e-3$, $1e-4$], yielding 9 hyperparameter combinations per MIL model. This ensured a fair evaluation after optimizing each method’s settings. The MIL models were trained for 50 epochs with a 5-epoch warmup and cosine annealing learning rate schedule. To select the best checkpoint for each representation, we split the training sets into 75% training and 25% validation partitions. The checkpoints with the highest validation set performance were chosen for final evaluation on the held-out test set (Table 2 and 3).

4.3. Ablation Study

To investigate the influence of the core components of our CluSiam model, we conducted an ablation study on the Camelyon16 dataset. We created and trained several model variants with different combinations of the invariance loss term, cluster loss term, and Gumbel noise to evaluate their impacts.

The stop gradient operation is crucial to prevent our model from collapsing, similar to its role in SimSiam. Additionally, incorporating Gumbel noise when assigning clusters is critical. Without this noise, the cluster assigner will collapse early in training, assigning all samples to one cluster. This collapsed state becomes equivalent to SimSiam training, as a single cluster removes the off-diagonal elements required to calculate the cluster loss $\mathcal{L}_{\text{cluster}}$. The gradient from $\mathcal{L}_{\text{cluster}}$ thus becomes zero in this situation. However, by introducing randomness in cluster assignment, the noise prevents updates from concentrating solely on the most probable cluster dimensions, thereby preventing early collapse. Together, the stop gradient and noise allow our model to escape these trivial single-cluster solutions, enabling effective joint optimization of \mathcal{L}_{inv} and $\mathcal{L}_{\text{cluster}}$. As




Figure 3: Cluster visualization for the unseen test set. Uncolored regions were filtered out during the preprocessing stage.

evidenced in Table 4, CluSiam outperforms the other models. We observed a collapse in the joint model’s cluster assigner, with one dimension consistently dominating the others. This is likely because only the largest dimension receives updates, continually amplifying its magnitude and leading to a single dominant cluster. The joint model was expected to mirror SimSiam’s performance, as its training dynamics should be identical to SimSiam after the assigner’s collapse. However, its performance was inferior compared to SimSiam. This discrepancy may originate from the cluster assigner offering suboptimal initialization for the subsequent pure SimSiam training before its collapse.

In the aftermath of our ablation studies on the new modules introduced to SimSiam, we delved into investigating the pivotal roles of original SimSiam components in ensuring training stability. A prime area of focus was the scale of inputs to the cluster assigner, as it might significantly influence this stability. The SimSiam projector, which interleaves \mathcal{BN} layers between its linear layers and concludes with a \mathcal{BN} layer, could be foundational for the clustering module’s effectiveness. To empirically assess the role of

controlled input scaling, we devised an experiment on two different projectors. The first projector, originating from SimSiam, concludes with a \mathcal{BN} layer. The second projector, from BYOL, ends with a linear layer. We began by replacing the SimSiam-style projector with the BYOL variant in our CluSiam model, leading to the creation of the CluBYOL model with the cluster assigner module integrated into the BYOL architecture. The CluBYOL was initially trained using the BYOL-style projector. Subsequently, we utilized the SimSiam-style projector and undertook another training round for the CluBYOL model. Notably, both situations using the BYOL-style projector resulted in collapse, akin to the joint model depicted in Table 4, with a single cluster predominantly emerging. As highlighted in Table 5, the SimSiam-style projector, characterized by its concluding with a \mathcal{BN} layer, is instrumental in preventing such collapses.

To further analyze the behavior of the clustering module, we examined the impact of the exploration space size, denoted as K . When Gumbel noise is introduced, some samples exhibit the highest probability of remaining in the cluster with the highest output value, while also possessing a

Agg.	Rep.	Acc.	AUC	
			Neg.	Pos.
Max	Supervised	0.628	0.421	0.501
	SimCLR	0.860	0.346	0.951
	SwAV	0.853	0.517	0.845
	PCL	0.496	0.370	0.510
	Barlow.	0.868	0.407	0.941
	BYOL	0.659	0.455	0.834
	SimSiam	0.690	0.316	0.680
	CluSiam	0.884	0.453	0.952
DSMIL	Supervised	0.651	0.635	0.635
	SimCLR	0.822	0.845	0.874
	SwAV	0.876	0.866	0.859
	PCL	0.488	0.535	0.496
	Barlow.	0.860	0.873	0.945
	BYOL	0.558	0.501	0.586
	SimSiam	0.721	0.656	0.680
	CluSiam	0.907	0.945	0.952

Table 2: Results on Camelyon16 dataset. The magnification level is $0.5 \mu\text{m}/\text{pixel}$. All the representations were trained using a batch size of 512 and the ResNet18 architecture.

Agg.	Rep.	Acc.	Neg.	AUC	
				Neo.	Pos.
Max	Supervised	0.359	0.313	0.562	0.494
	SimCLR	0.462	0.565	0.549	0.720
	SwAV	0.462	0.497	0.451	0.726
	PCL	0.692	0.556	0.935	0.731
	Barlow.	0.538	0.438	0.509	0.843
	BYOL	0.462	0.314	0.719	0.589
	SimSiam	0.359	0.598	0.531	0.694
	CluSiam	0.641	0.598	0.966	0.851
DSMIL	Supervised	0.356	0.296	0.617	0.529
	SimCLR	0.718	0.497	0.904	0.700
	SwAV	0.538	0.669	0.840	0.726
	PCL	0.744	0.710	0.969	0.797
	Barlow.	0.615	0.672	0.957	0.831
	BYOL	0.564	0.527	0.981	0.697
	SimSiam	0.590	0.686	0.910	0.803
	CluSiam	0.769	0.754	0.985	0.883

Table 3: Results on the Pancreatic Cancer dataset. The magnification level is $0.5 \mu\text{m}/\text{pixel}$. All the representations were trained using a batch size of 512 and the ResNet18 architecture.

high probability of transitioning to a nearby or similar cluster. Some samples that are difficult to differentiate might be allocated with near-equal probabilities across multiple centroids. A larger value of K results in more refined clustering. Conversely, when K is small, achieving definitive clustering becomes difficult. For instance, with K set to 3, the

Model	\mathcal{L}_{inv}	$\mathcal{L}_{cluster}$	Noise	Acc.	AUC	
					Neg.	Pos.
SimSiam	✓	-	-	0.721	0.656	0.680
Cluster	-	✓	✓	0.426	0.502	0.627
Joint	✓	✓	-	0.667	0.538	0.563
CluSiam	✓	✓	✓	0.907	0.945	0.952

Table 4: Ablation study of CluSiam components on Camelyon16 using DSMIL aggregation.

Model	Projector Ending with \mathcal{BN}	Acc.	AUC	
			Neg.	Pos.
CluSiam	-	0.643	0.556	0.591
	✓	0.907	0.945	0.952
CluBYOL	-	0.627	0.656	0.658
	✓	0.923	0.947	0.975


Table 5: Results on Camelyon16 using DSMIL aggregation.

cluster assigner can allocate some hard-to-distinguish samples into a third cluster, roughly equidistant from the first two. Yet, with K set to 2, the assigner can only place samples into one of the two clusters. Importantly, when K is 1, the model essentially becomes SimSiam since its loss function and backpropagation are equivalent to those in SimSiam. In this scenario, the cluster assigner lacks the flexibility to differentiate between samples by assigning them to different clusters.


In our experiments, we assessed the influence of K on model behavior by training two models with exploration spaces of $K = 10$ and $K = 100$. Using a Top-1 KNN classifier’s F1 score for patch-level performance evaluation, both models displayed comparable classification performance, as depicted in Figure 4. Despite this similarity in classification, their clustering behaviors were distinct. The model with $K = 10$ exhibited a larger number of clusters and greater fluctuations in both cluster counts and the Rand Index. In contrast, the larger exploration space of $K = 100$ allowed the assigner to stabilize on more definitive assignments faster. This difference underlines the influence of K in CluSiam’s clustering. Specifically, both the number of clusters and the Rand Index fluctuate more with $K=10$. This limited exploration space prevents highly fine-grained cluster assignments. In contrast, a larger exploration space of $K=100$ provides more granularity for refined clustering actions, enabling the assigner to stabilize on more definitive assignments rapidly. The cluster count also becomes more consistent with $K=100$. These observations highlight the impact of K on CluSiam’s clustering dynamics.

5. Conclusion

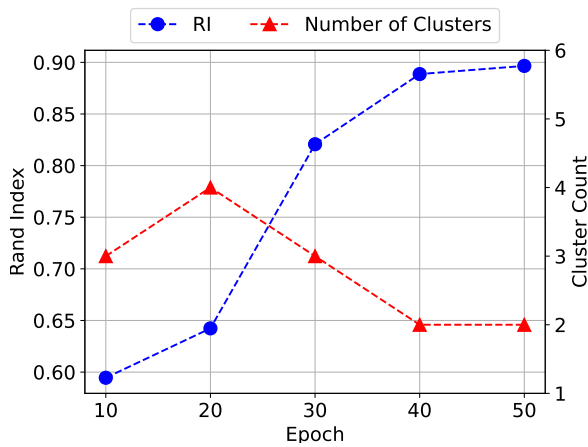
In this paper, we introduce CluSiam, a SSL technique that integrates cluster constraints to enhance representation



(a) Patch-level F1 scores using the Top-1 KNN classifier.



(b) WSI-level clustering patterns with exploration space $K=10$.



(c) WSI-level clustering patterns with exploration space $K=100$.

Figure 4: Comparison of models with different exploration spaces K on Camelyon16.

learning for histopathology images. By subtly pushing apart inter-cluster instances while aligning intra-cluster views, CluSiam balances similarity and dissimilarity. It demonstrates substantial improvements in downstream classification and clustering tasks compared to baseline methods. Additionally, CluSiam provides an efficient way to analyze histopathology images without requiring manual annotations.

6. Acknowledgements

This research was supported in part by grants from the US National Library of Medicine (R01LM012837 and R01LM013833) and the US National Cancer Institute (R01CA249758).

References

- [1] Murtaza Ashraf, Willmer Rafell Quiñones Robles, Mujin Kim, Young Sin Ko, and Mun Yong Yi. A loss-based patch label denoising method for improving whole-slide image analysis using a convolutional neural network. *Scientific Reports*, 12(1):1392, 2022.
- [2] Adrien Bardes, Jean Ponce, and Yann LeCun. Vircreg: Variance-invariance-covariance regularization for self-supervised learning. *arXiv preprint arXiv:2105.04906*, 2021.
- [3] Gabriele Campanella, Matthew G Hanna, Luke Geneslaw, Allen Mirafior, Vitor Werneck Krauss Silva, Klaus J Busam, Edi Brogi, Victor E Reuter, David S Klimstra, and Thomas J Fuchs. Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nature medicine*, 25(8):1301–1309, 2019.
- [4] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European conference on computer vision (ECCV)*, pages 132–149, 2018.
- [5] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in neural information processing systems*, 33:9912–9924, 2020.
- [6] 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.
- [7] Jacob Carse, Frank Carey, and Stephen McKenna. Unsupervised representation learning from pathology images with multi-directional contrastive predictive coding. In *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, pages 1254–1258. IEEE, 2021.
- [8] 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, 2022.
- [9] Chi-Long Chen, Chi-Chung Chen, Wei-Hsiang Yu, Szu-Hua Chen, Yu-Chan Chang, Tai-I Hsu, Michael Hsiao, Chao-Yuan Yeh, and Cheng-Yu Chen. An annotation-free whole-slide training approach to pathological classification of lung cancer types using deep learning. *Nature communications*, 12(1):1193, 2021.
- [10] 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.
- [11] 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.
- [12] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15750–15758, 2021.
- [13] Hsien-Tzu Cheng, Chun-Fu Yeh, Po-Chen Kuo, Andy Wei, Keng-Chi Liu, Mong-Chi Ko, Kuan-Hua Chao, Yu-Ching Peng, and Tyng-Luh Liu. Self-similarity student for partial label histopathology image segmentation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXV 16*, pages 117–132. Springer, 2020.
- [14] Marina D’Amato, Przemysław Szostak, and Benjamin Torben-Nielsen. A comparison between single-and multi-scale approaches for classification of histopathology images. *Frontiers in Public Health*, 10, 2022.
- [15] Joseph DiPalma, Arief A Suriawinata, Laura J Tafe, Lorenzo Torresani, and Saeed Hassanpour. Resolution-based distillation for efficient histology image classification. *Artificial Intelligence in Medicine*, 119:102136, 2021.
- [16] Joseph DiPalma, Lorenzo Torresani, and Saeed Hassanpour. Histoperm: A permutation-based view generation approach for improving histopathologic feature representation learning. *Journal of Pathology Informatics*, page 100320, 2023.
- [17] Debidatta Dwivedi, Yusuf Aytar, Jonathan Tompson, Pierre Sermanet, and Andrew Zisserman. With a little help from my friends: Nearest-neighbor contrastive learning of visual representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9588–9597, 2021.
- [18] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020.
- [19] 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.
- [20] 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.
- [21] Le Hou, Dimitris Samaras, Tahsin M Kurc, Yi Gao, James E Davis, and Joel H Saltz. Patch-based convolutional neural network for whole slide tissue image classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2424–2433, 2016.
- [22] Maximilian Ilse, Jakub Tomczak, and Max Welling. Attention-based deep multiple instance learning. In *International conference on machine learning*, pages 2127–2136. PMLR, 2018.
- [23] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.
- [24] Andrew Janowczyk and Anant Madabhushi. Deep learning for digital pathology image analysis: A comprehensive tutorial with selected use cases. *Journal of pathology informatics*, 7(1):29, 2016.

- [25] Yiren Jian, Chongyang Gao, and Soroush Vosoughi. Non-linguistic supervision for contrastive learning of sentence embeddings. *Advances in Neural Information Processing Systems*, 35:35533–35548, 2022.
- [26] Shuai Jiang, Liesbeth Hondelink, Arief A. Suriawinata, and Saeed Hassanpour. Masked pre-training of transformers for histology image analysis, 2023.
- [27] 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.
- [28] Han Le, Dimitris Samaras, Tahsin Kurc, Rajarsi Gupta, Kenneth Shroyer, and Joel Saltz. Pancreatic cancer detection in whole slide images using noisy label annotations. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part I 22*, pages 541–549. Springer, 2019.
- [29] Marvin Lrousseau, Maria Vakalopoulou, Marion Classe, Julien Adam, Enzo Battistella, Alexandre Carré, Théo Estienne, Théophraste Henry, Eric Deutsch, and Nikos Paragios. Weakly supervised multiple instance learning histopathological tumor segmentation. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part V 23*, pages 470–479. Springer, 2020.
- [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 *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14313–14323, Nashville, TN, USA, June 2021. IEEE.
- [31] Junnan Li, Pan Zhou, Caiming Xiong, and Steven Hoi. Prototypical contrastive learning of unsupervised representations. In *International Conference on Learning Representations*.
- [32] Geert Litjens, Peter Bandi, Babak Ehteshami Bejnordi, Oscar Geessink, Maschenka Balkenhol, Peter Bult, Altuna Halilovic, Meyke Hermesen, Rob van de Loo, Rob Vogels, et al. 1399 h&e-stained sentinel lymph node sections of breast cancer patients: the camelyon dataset. *GigaScience*, 7(6):giy065, 2018.
- [33] Yun Liu, Krishna Gadepalli, Mohammad Norouzi, George E Dahl, Timo Kohlberger, Aleksey Boyko, Subhashini Venugopalan, Aleksei Timofeev, Philip Q Nelson, Greg S Corrado, et al. Detecting cancer metastases on gigapixel pathology images. *arXiv preprint arXiv:1703.02442*, 2017.
- [34] Caner Mercan, Selim Aksoy, Ezgi Mercan, Linda G Shapiro, Donald L Weaver, and Joann G Elmore. Multi-instance multi-label learning for multi-class classification of whole slide breast histopathology images. *IEEE transactions on medical imaging*, 37(1):316–325, 2017.
- [35] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [36] Naofumi Tomita, Behnaz Abdollahi, Jason Wei, Bing Ren, Arief Suriawinata, and Saeed Hassanpour. Attention-Based Deep Neural Networks for Detection of Cancerous and Precancerous Esophagus Tissue on Histopathological Slides. *JAMA Network Open*, 2(11):e1914645–e1914645, 11 2019.
- [37] Naofumi Tomita, Laura J Tafe, Arief A Suriawinata, Gregory J Tsongalis, Mustafa Nasir-Moin, Konstantin Dragnev, and Saeed Hassanpour. Predicting oncogene mutations of lung cancer using deep learning and histopathologic features on whole-slide images. *Translational Oncology*, 24:101494, 2022.
- [38] Wouter Van Gansbeke, Simon Vandenhende, Stamatios Georgoulis, Marc Proesmans, and Luc Van Gool. Scan: Learning to classify images without labels. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part X*, pages 268–285. Springer, 2020.
- [39] Noorul Wahab, Islam M Miligy, Katherine Dodd, Harvir Sahota, Michael Toss, Wenqi Lu, Mostafa Jahanifar, Mohsin Bilal, Simon Graham, Young Park, et al. Semantic annotation for computational pathology: Multidisciplinary experience and best practice recommendations. *The Journal of Pathology: Clinical Research*, 8(2):116–128, 2022.
- [40] Dayong Wang, Aditya Khosla, Rishab Gargeya, Humayun Irshad, and Andrew H Beck. Deep learning for identifying metastatic breast cancer. *arXiv preprint arXiv:1606.05718*, 2016.
- [41] Xiyue Wang, Yuexi Du, Sen Yang, Jun Zhang, Minghui Wang, Jing Zhang, Wei Yang, Junzhou Huang, and Xiao Han. Retccl: Clustering-guided contrastive learning for whole-slide image retrieval. *Medical Image Analysis*, 83:102645, 2023.
- [42] Jason W Wei, Laura J Tafe, Yevgeniy A Linnik, Louis J Vaickus, Naofumi Tomita, and Saeed Hassanpour. Pathologist-level classification of histologic patterns on resected lung adenocarcinoma slides with deep neural networks. *Scientific reports*, 9(1):3358, 2019.
- [43] Weiyi Wu, Xiaoying Liu, Robert B Hamilton, Arief A Suriawinata, and Saeed Hassanpour. Graph convolutional neural networks for histologic classification of pancreatic cancer. *Archives of Pathology & Laboratory Medicine*, 2023.
- [44] Jiawei Yang, Hanbo Chen, Yuan Liang, Junzhou Huang, Lei He, and Jianhua Yao. Concl: Concept contrastive learning for dense prediction pre-training in pathology images. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXI*, pages 523–539. Springer, 2022.
- [45] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *International Conference on Machine Learning*, pages 12310–12320. PMLR, 2021.
- [46] Yu Zhao, Fan Yang, Yuqi Fang, Hailing Liu, Niyun Zhou, Jun Zhang, Jiarui Sun, Sen Yang, Bjoern Menze, Xijuan Fan, et al. Predicting lymph node metastasis using histopathological images based on multiple instance learning with deep graph convolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4837–4846, 2020.