



Enhancing Prostate Cancer Subtyping: A Persistent Homology Approach in Multiple Instance Learning

Ahmad Obeid,* Sajid Javed, Jorge Dias, Ibrahim Abe M. Elfadel and Naoufel Werghi

Abstract

Prostate pathology analysis has been handled remarkably well by deep learning models. However, deep learning models in histopathology still face challenges in effectively modeling fine-grained, object-level features that are critical for a robust cancer assessment. Topological Data Analysis (TDA) has shown promise in addressing these issues but remains underexplored, particularly for whole-slide pathology applications. This is manifested in the nonexistent application of TDA on the slide level; an issue that significantly undermines their practical usability, and the non-utilization of the multi-magnification nature of Whole-slide Images (WSIs). Furthermore, existing studies are limited to small-scale datasets, thus challenging the validity of TDA in histopathology. In this work, we address these gaps by introducing Persistent Homology in Multiple Instance Learning (PMIL), the first adaptable TDA-based module within the MIL framework. We further propose the Cubic version cPMIL, utilizing the magnification factor in pathology images for filtration and unlocking an improved object-level modeling capability. We validate our approach on two prostate subtyping datasets, comparing against multiple state-of-the-art methods. Our proposed modules are developed offline, can be seamlessly plugged in any MIL pipeline, do not incur any additional expert annotation, operational within clinical limitations, user-friendly, and open-source. We provide a detailed guide for implementing the modules at github.com/ahmadobeid/PMIL.

Keywords: Deep learning; Histopathology; Image classification; Medical imaging; Topology.

Received: 08 June 2025; Revised: 26 October 2025; Accepted: 19 November 2025

Article type: Research article.

1. Introduction

Whole-slide images (WSI) have emerged as a cornerstone in computational pathology (CPath), where deep learning (DL) tools increasingly enhance or even replace traditional tasks such as mitosis detection, biopsy classification and tissue segmentation. Additionally, DL methods show promise in minimizing inter-observer variability and intra-observer inconsistency, while driving novel treatment strategies. This advancement made the clinical adoption of DL tools within reach, especially given the concurrent progress in medical AI applications, including diagnostic support through facial expression analysis,^[1] rehabilitation systems using EMG signals,^[2] and ultrasound-based disease classification,^[3] in addition to other clinical domains.^[4]

The preparation of a histopathological slide consists of several phases, including the collection, staining, storage, and annotation stages. Oftentimes, artifacts and noise may occur

in any of the stages. Compounded by sparse or inconsistent labels, the analysis of the heterogeneous cancer landscape, consisting of varied tissue structures, cellular overlap, irregular shapes, inflammation, necrosis, and inconsistent staining becomes a highly complex task. As such, advanced methods and algorithms are required, particularly as pertains to object-level modelling of nuclear and glandular structures in the histopathology image. Successfully achieving the latter while mitigating the effect of noise can lead to better decision making and interpretability.

While morphological and chromatic features are usually successfully modelled through the convolution/attention operations in the RGB domain, the topology signature of the image offers irreplaceable insights that aid in image analysis and adds an extra layer of robustness against noisy elements in the RGB image. TDA offers a powerful method to analyse the shape and structure of data at both local and global scales. By modelling the data as a point cloud, TDA monitors the emergence (birth), fading (death), and persistence of topological features, such as connected components, loops, and voids, as a certain filtration parameter progresses. This

College of Computing and Mathematical Sciences, Khalifa University,

Abu Dhabi, 127788, UAE

*Email: ahmad.obaid@ku.ac.ae (A. Obeid)

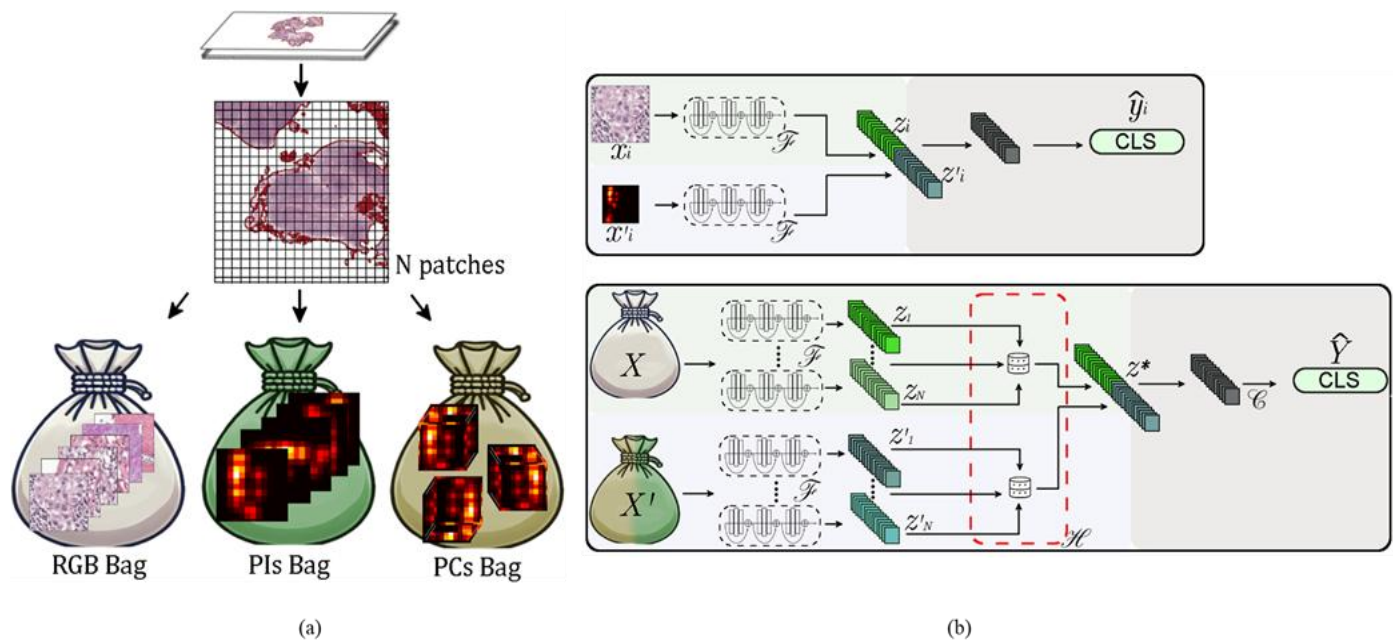


Fig. 1: TDA and MIL in histopathology. Patches are extracted from WSIs, from which we can create Bags of RGB images, Persistence Images (PIs), or Persistence Cubes (PCs) (a), both of them examples of Persistence Bags (PBs). The classical application of TDA in CPath (b, upper) is limited to the patches. Extending TDA to the full slide utilizes the MIL paradigm (b, lower), where either bags of Persistence Images or bags of Persistence Cubes are analysed. An aggregation network is used for RGB and PIs/PCs separately, after which, they are combined at the feature level.

persistence effectively identifies meaningful patterns that are pertinent to prostate cancer subtyping. Secondly, these patterns are distinguished from noise and artifacts, as the modelling provides a natural filtering procedure of nonpersistent patterns. Finally, the extracted information is represented using precise tools like the Persistence Barcode or Persistence Diagram (PD), enabling robust, noise-resistant modelling of pathology images at the object level, including glands and nuclei.

While TDA offers a promising avenue for capturing object-level features in histopathology, its adoption remains hindered by several limitations. Existing approaches are confined to patch-level analysis and rely on detailed instance annotations. These settings diverge from clinical workflows and limit scalability. The PMIL module addresses these issues by introducing persistence-based topological features into the Multiple Instance Learning (MIL) framework, enabling full-slide analysis using only slide-level labels. By creating Persistence Bags (PBs) of topological summaries (e.g., persistence images), PMIL captures structural patterns such as nuclear clustering and gland formation, and integrates them with standard RGB features in a modular, annotation-free manner. The cPMIL module further extends this by incorporating the magnification factor as a second filtration parameter. This allows cPMIL to extract scale-aware topological signatures through persistence cubes, capturing both fine-grained and coarse tissue structures across different magnification levels—thus mimicking the multi-scale reasoning used by pathologists during diagnosis.

1.1 Related works

Our work lies at the intersection of three main domains: Topological Data Analytics, Multiple Instance Learning in histopathology, and multi-parameter filtration. We will delineate the relevant works to better contextualize and highlight our contributions.

A. Topological data analysis in CPath

Incorporating topological features into DL models offers significant advantages, extending beyond the capabilities of convolutional and attention-based methods to capture and preserve categorically distinct features. Recently, TDA-based approaches have shown promising results in various machine learning (ML) and computer vision applications, especially histopathology analysis.^[5,6]

For example, TDA has been used to differentiate glandular structures between tumor and non-tumor regions.^[7] Another study employed persistence homology profiles from slide patches within a K-Nearest Neighbour (KNN) model for classification.^[8] Similarly, modelling glandular topology with TDA enhanced prostate cancer classification, outperforming the traditional Gleason grading system.^[9,10] The Cell Community Forest (CCF)^[11] integrates topological, appearance, and microenvironmental features, leveraging a graph neural network to combine features across different levels for a comprehensive image-level diagnosis. Additionally, PHG-Net^[12] uses pixel intensity for filtration and encodes each point in the PD as a high-dimensional feature through multi-layer perceptrons (MLPs).

Further advances, such as Persistence Images (PI) and Persistence Landscapes (PL),^[13,14] which transform PD into vectorized representations, have improved the compatibility of TDA with DL models. Faster TDA computation^[7] has also increased its applicability in research.

Despite these strengths, the potential for TDA in pathology remains underexplored, with existing applications limited primarily to patch-level analysis on small datasets. In turn, this makes the application of TDA in histopathology contingent on the existence of patch labels; a highly demanding condition which hinders its adoption in clinical practice. In this work, our aim is to align with data-efficient approaches such as the Multiple Instance Learning (MIL) framework, which requires only slide-level labels. We integrate TDA into slide-level classification and validate its utility in CPath through experiments on challenging large-scale datasets.

B. Multiple instance learning in CPath

Slide-level classification in histopathology often uses the MIL framework, which we adopt to implement TDA at the slide level. MIL methods are well-suited for handling slide-level labels, offering a resource-efficient alternative to obtaining detailed, fine-grained annotations. Attention-based MIL approaches, such as the Attention-Based Multiple Instance Learning (ABMIL) framework^[15] improve interpretability by pinpointing key instances for classification, surpassing traditional methods like max-pooling and mean-pooling.^[16]

The Clustering-Constrained Attention MIL (CLAM) model^[17] extends ABMIL by introducing separate attention pathways for different classes. Similarly, the Dual-Stream MIL (DSMIL) model^[18] employs a dual-strategy mechanism to refine attention relationships, emphasizing critical instances. The double-tier feature distillation MIL (DTFD-MIL) model^[19] further builds on ABMIL by integrating grad-CAM into the attention pathway, creating smaller “pseudo bags” from large data collections (“bags”) to extract features at both the pseudo and parent bag levels. Finally, the Attention-Challenging (ACMIL) model^[20] aims to mitigate the over-concentration on a small subset of discriminative patches while ignoring a vast collection of other patches - a phenomenon observed in most attention-based MIL methods, and may be a source of overfitting. The authors utilize a multi-branch attention approach to accommodate the different patterns observed uniquely by different subsets of the slide. Additionally, the authors propose to randomly mask a subset of top-K discriminative instances, and assign their attention score to other, under-explored instances.

Although these MIL approaches have significantly advanced slide-level pathology analysis and improved interpretability, their reliance on convolution and attention mechanisms in the RGB domain often limits their ability to achieve fine-grained object-level modelling. This focus on data efficiency comes at the expense of capturing the intricate structural and spatial relationships within histopathology images. The integration of topological representations, which

can offer a richer semantic understanding, remains largely unexplored. TDA, with its straightforward implementation and seamless compatibility with existing workflows, presents an untapped opportunity to address this gap. By enabling a deeper and more robust modelling of histopathology objects—such as nuclei neighbourhoods and glandular structures—TDA has the potential to elevate the MIL paradigm to a new level of precision and insight.

C. Multiparameter filtration

Existing TDA approaches in CPath are constrained not only due to their restriction to patch-level analysis but also their omission of the inherent multi-magnification information available in most histopathology datasets—a factor routinely utilized in clinical practice. To address this, we introduce a multi-parameter TDA approach as a means to exploit the multi-magnification potential of histopathology images. The theory of multi-parameter Persistent Homology (MPH) has garnered considerable attention in the literature,^[21-23] with notable advancements such as the work of Vipond *et al.*^[24] in CPath. Their method enhances basic TDA by incorporating an additional filtration parameter, co-density, to improve noise robustness. Like other approaches, the analysis of^[24] begins with a point cloud of cells and is aligned with our own use of multiparameter TDA in CPath. However, it overlooks the critical multi-magnification factor inherently present in histopathology datasets.

The contributions of this paper can be summarized as follows:

1. We propose PMIL, the first MIL-based CPath module that integrates Persistence Homology at the whole-slide level
2. We propose cPMIL, a cubic version of PMIL which utilizes the multi-magnification factor in pathology images.
3. We implement PMIL and cPMIL in a modular fashion, allowing their deployment as a simple add-on to MIL approaches
4. We run experiments on a large-scale pathology dataset and benchmark PMIL and cPMIL against several state-of-the-art methods, illustrating the ease and benefits of integrating the PMIL and cPMIL modules within Multiple-Instance Learning frameworks.

2. Methodology

In this section, we describe the general steps followed to establish the TDA pipeline, and integrate it in the MIL paradigm. Furthermore, we describe our novel magnification-based TDA method.

2.1 Extracting persistent homology features

To analyse a histopathology image from a topology perspective is to quantify, analyse, and encode the geometric properties that are preserved under continuous deformations such as stretching and bending. Specifically, we are concerned with nuclear neighbourhoods and glandular structures; two highly beneficial biomarkers in cancer diagnosis, and

especially prostate subtyping. To that end, we focus on homology features, particularly 0^{th} -level homology (H_0) and 1^{st} -level homology (H_1). H_0 measures the connectedness of a point cloud, while H_1 captures loop-like structures within the data. By establishing a point-cloud that well represents the tumor microenvironment, we are able to condense a large amount of useful data into a precise and small-scale representation.

We begin by extracting nuclear objects, and generate a point cloud reflecting the nuclei's spatial placement within the microenvironment. For this step, a pre-trained nuclei detector can be used such as the HoVer-Net model,^[25] as well as simple thresholding and morphological operations. In this work, we utilize the latter over each patch to extract its nuclear content. We follow the work of^[7] which involves a standardized pipeline: patches are first color-normalized using the Reinhard method,^[26] followed by color deconvolution via Macenko's approach to isolate the nuclear stain^[27] Minimum cross-entropy thresholding is then applied to segment the nuclear foreground. After normalizing the stain and segmenting the nuclear foreground, nuclei centroids are detected using a scale-adaptive Laplacian-of-Gaussian (LoG) filter, implemented as a fast Difference-of-Gaussian (DoG) approximation. This multiscale blob detection method effectively identifies nucleus-like structures, and is implemented in the open-source HistomicsTK toolkit. The arrangement of nuclei, their clutter, shape, and distribution are key elements in the diagnosis of a histopathology image and are encoded effectively using the H_0 feature. Additionally, glands are identified as a circular arrangement of lymphocytes, and can be captured and analysed through the H_1 feature. As

such, the established point-cloud data is transformed into an instrumental representation of key biomarkers in the image that enhance the accuracy and sensitivity of cancer diagnosis. This process is illustrated in Fig. 2 with two pathology image examples.

The process of constructing a topological signature from a point cloud begins by placing geometric balls with an initial radius r around each nucleus. As r gradually increases, nearby nuclei merge, forming local neighbourhoods known as a simplicial complex ζ_i . This expansion continues until all nuclei collapse into a single connected component as $r \rightarrow \infty$. Each step in this filtration is recorded, resulting in the PD shown in Fig. 2(c). This structured growth, where a parameter-controlled complex evolves through a nested sequence ($\zeta_0 = \emptyset \subset \zeta_1 \subset \zeta_2 \dots \subset \zeta_n =$), is referred to as a filtration process. The simplicial complex derived from a point cloud (nuclei in our case), is known as a Vietoris-Rips complex. By analysing the profiles of H_0 and H_1 , we effectively characterize the spatial organization of nuclei and the structural distribution of glandular formations, respectively.

A comparison of the PDs from two histopathology images illustrates the significance of topological features as shown in Fig. 2. In the upper image, which contains numerous glandular structures, many persistent H_1 features emerge early and remain for an extended duration, while these features are absent in the lower image. Conversely, the lower image exhibits a larger number of connected components (H_0) that persist longer. At the final radius, the upper image condenses into a single connected component, whereas the lower image retains multiple components at the same radius. These differences highlight TDA's effectiveness in object-level

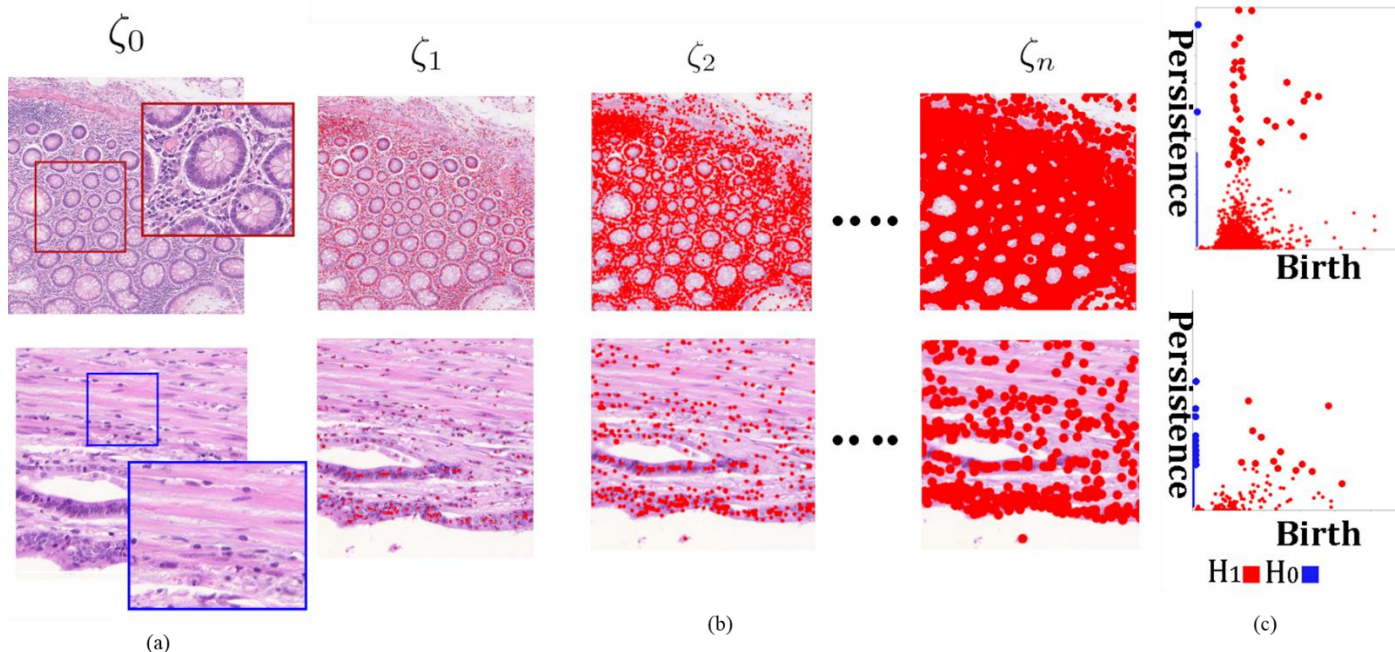


Fig. 2: Persistence Diagram Extraction in CPath. Two pathology images are shown in (a). The upper row has more glandular objects than the lower row, which is less cluttered. Filtration using balls centred on each nucleus is shown in (b). The resulting Persistence Diagram for each image is shown in (c).

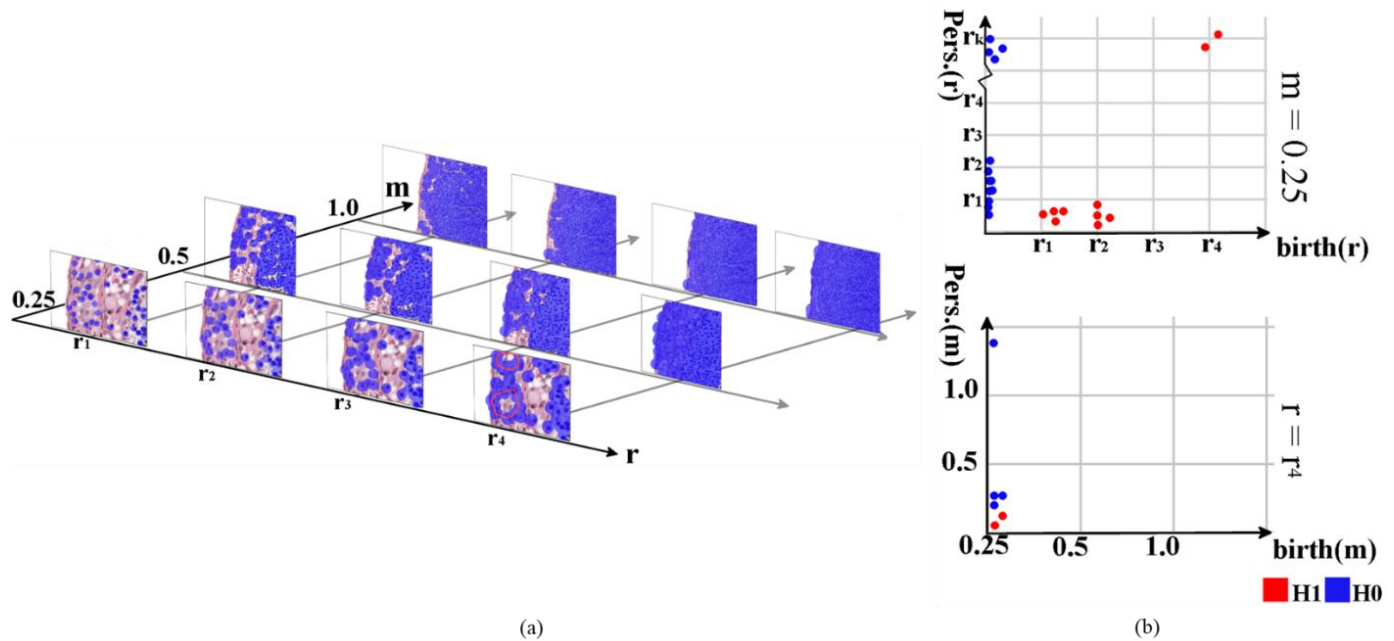


Fig. 3: The cPMIL module. Both the radius r and the magnification factor m are suitable filtration parameters as shown in (a). At radius r_4 , two complete glands are discovered and persist for a long filtration range as revealed in (b, upper plot). In (b), the persistence diagrams are generated by fixing one parameter and filtering according to the other.

modelling, reinforcing findings from previous studies.^[9]

2.2 Multiple instance learning

With the rapid advancement in sensing technology, modern WSIs typically span extremely high resolutions and feature rich information content, which raises computational and memory challenges in most typical DL models. Consequently, DL pipelines often start by dividing the WSI into numerous smaller patches as illustrated in the first step in Fig. 1. In the MIL setting, the WSI is treated as a collection (“bag”), with each patch representing an element (“instance”) within the bag. Consider a WSI divided into N patches $x_i \in X$. Let $Y \in \{0,1\}$ represent the label for the WSI X , and y_i denote the label for each patch x_i .

Due to resource limitations, it is typical that a pathologist would designate a label for the entire slide, without providing labels for the individual instances. In fact, this is where past attempts to integrate TDA in CPath diverge from clinical practice, that is, assuming accessibility to individual labels y_i . In contrast, the MIL setting allows access to Y only, making it more realistic and representative of clinical practice, yet more challenging to employ.

The fundamental assumption of MIL posits that if $Y = 0$, then it follows that $\sum y_i = 0$, and if $Y = 1$ then $\sum y_i > 0$. MIL is typically operated according to:

$$z_i = \mathcal{F}(x_i), Z = \mathcal{H}(z_i), \hat{Y} = \mathcal{C}(Z) \tag{1}$$

where \mathcal{F} represents an instance-level feature extractor, \mathcal{H} denotes an aggregation network which assembles Z from the individual z_i instances, and \mathcal{C} is a classifier head. Oftentimes, networks \mathcal{H} and \mathcal{C} are combined.

2.3 Persistence images in MIL

The Persistence Images in MIL (PMIL) module proposed in this work extends the MIL framework beyond the RGB domain. Fig. 1 illustrates the distinction between PMIL and the conventional approach to using TDA in histopathology. Similar to constructing bags of patches, the PMIL module creates bags of PIs. We refer to these as Persistence Bags (PBs) X' . Each instance $x'_i \in X'$ is associated with an RGB patch x_i , and they share the same label Y . The PMIL module analyses the PBs independently from their RGB counterparts, and a later fusion layer combines the two in a unified representation which integrates topological information with the typical chromatic/morphological features via the RGB-based operations.

Several options exist to fuse the features of the PMIL module with the main stream, including feature fusion, decision-level fusion, and regularization via combined loss terms. In this work, we apply feature-level fusion using self-attention over the combined feature vector $Z^* = Z \# Z'$, where Z and Z' are the aggregated representations in Eq. (1) of the RGB and persistence images, respectively, and $\#$ denotes concatenation. Thus, Eq. (1) is modified to $\hat{Y} = \mathcal{C}(Z^*)$.

2.4 Magnification as a filtration parameter

Most existing approaches follow the basic procedure of extracting topology features described in Section 2.1 **Extracting persistent homology features**, where the filtration parameter is the radius r of balls centred around each point in the cloud. In this work, we propose the utilization of an inherent complementary varying parameter, the magnification factor m , which governs the scale at which

image patches are viewed and analysed. In particular, texture information, local connections including circular arrangements in glands, and other high-frequency phenomena are pronounced at high magnification, while global connections and other large-scale phenomena such as tumor-infiltrating lymphocytes (TILs), stromal distributions, and large-scale tissue architecture are filtered out. Conversely, at low magnification, global structural features and large-scale patterns dominate, while finer, high-frequency details are effectively filtered out. This duality in representation underscores the role of magnification as a filtration parameter, whereby varying m for the same radius r creates or eliminates significant topological biomarkers that are scale-dependent. This is demonstrated visually in Fig. 3(a).

The persistence diagrams in Fig. 3(b) provides further insight. In the upper diagram, which represents the typical case, all initial connected components (H_0 objects), are born at early radii, converge gradually, and are terminated along the advancement of the radius parameter. Moreover, two persistent H_1 objects are born late into the filtration. On the complementary view, the lower diagram is generated by fixing the radius at r_4 , at which the two H_1 objects and the persistent H_0 objects are born, and decreasing the magnification factor. At the second stage of the filtration through m , the two H_1 components are eliminated, alluding to their relatively-small sizes, which agrees with the visual assessment of the corresponding glands in the image.

Concretely, consider the region depicted at the highest magnification $m = 0.25 \mu\text{m}/\text{pixel}$ (μpp) in Fig. 3(a). The standard extraction of topological features identifies the existence of two glands as can be seen at r_4 . It is also possible to approximate the connectedness of nuclei through the average of the standard H_0 persistence values. On the flip-side, by observing magnification m as a filtration parameter and holding r constant at r_4 , decreasing the magnification view immediately terminates the two glands, providing an indirect

yet intuitive estimate of gland sizes—a well-established biomarker with significant diagnostic value in histopathology.^[28-30] Similarly, this approach applies to the assessment of nuclear density: compact nuclear regions collapse into fewer connected components at higher magnifications, while sparse regions remain disconnected for longer. This additional layer of information enables a more nuanced characterization of histopathological features, linking gland size and nuclear clutter to clinically relevant metrics.

Moreover, considering the 2D case combining the variation of both parameters reveals additional benefits. In detail, observing the r -based persistence at multiple m values, in contrast to a single value, provides a more holistic characterization of the tissue environment and establishes M persistence diagrams, $PD(m), 1 \leq m \leq M$, where M represents the available number of magnification views in the dataset, and $PD(m)$ is the persistence diagram represented as a function of the parameter m . The basic implementation of TDA assumes $M = 1$, limiting the analysis to a single magnification *e.g.* $0.25 \mu\text{pp}$. For example, when analysing TILs at such a magnification, the connectedness of the point-cloud data representing infiltrating lymphocytes can be assessed through H_0 , revealing immune engagement and cell-to-cell interaction. At smaller magnifications, $m = 1.0 \mu\text{pp}$, the spread of the immune response becomes discernible, offering complementary insights. By modelling these interactions across different scales within the tissue environment, this multi-scale approach enables a more comprehensive diagnosis of the slide.

2.5 Filtering through magnification

To model the 2D interactions as described above, the direct method involves building a 3D Persistence Diagram, from which a full persistence cube is extracted. However, unlike the parameter r , the magnification factor m is not continuous, as most existing datasets provide a finite set of magnification

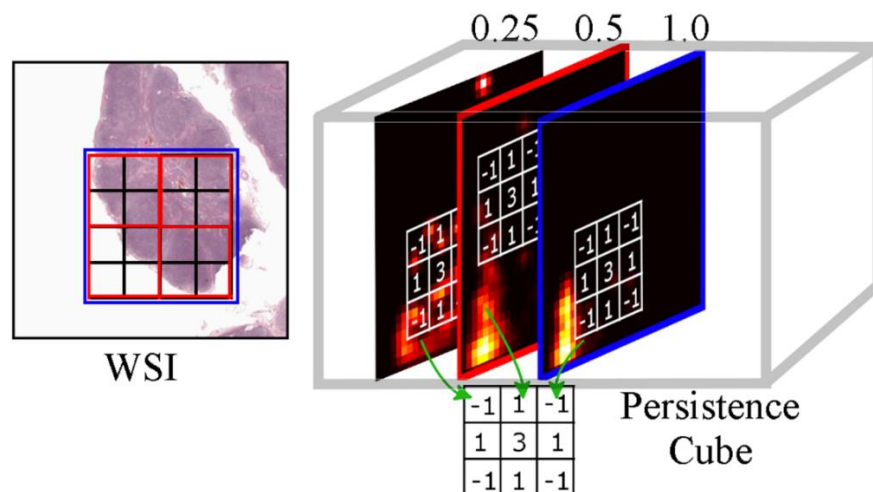


Fig. 4: Patching the WSI is done on several magnifications. In the left image, both $1.0 \mu\text{pp}$ (blue) and $0.5 \mu\text{pp}$ (red) are shared by several regions. On the right, the PC is formed out of the different views, where across-radius and across-magnification convolution are both shown.

priors--usually around three to six. Consequently, to simplify the implementation of the 2D approach while preserving both cross-magnification and cross-radius interactions, we construct a sliced cubic representation of each regional topology persistence, where each slice corresponds to a magnification factor. In the persistence cube (PC), the highest magnification patch is centred around the region of interest, while lower magnifications extend to other regions. This is illustrated in Fig. 4. In other words, all PIs, except that of the highest magnification, are shared across neighbouring persistence PCs. Next, a 3D convolution is applied to the PC, thus modelling intra-magnification connections associated with the filtration parameter r , and inter-magnification connections where the filtration parameter is m .

2.6 Optimizing performance, memory and computation

The proposed modules explore an additional source of information that assists in the classification of cancers from histopathology images. The added information is represented in a condensed manner, allowing the improvement to come at a modest computational/memory overhead. In particular, the resolution of a persistence bag instance, x'_i , is generally much smaller than x_i , making the storage of the additional bag and its processing inexpensive. This is especially pronounced with the cPMIL module. The bag of PCs does not only include an additional modality, but also expands to multiple magnifications. Higher magnification values may include far more instances representing a region of interest. In this case, the alternative, represented by creating RGB bags at higher magnification levels, would require a prohibitively expensive memory footprint. Worse still, the following cross-magnification attention/convolution will also be prohibitively expensive. In detail, we estimate that this naive alternative would incur up to $4 - 6 \times$ overhead in parameters, computation, and memory as illustrated in Table 1. On the other hand, with the PMIL and cPMIL modules, the overall number of trainable parameters and FLOPs remains low, increasing only by a factor of $1.66 \times$ over vanilla MIL. Using the topology domain, we encode much valuable object-level information of a WSI into a small-scale, easy-to-analyse packet.

A key advantage of our approach is the efficiency of computing persistence images for individual patches. Following,^[7] the generation of a single persistence image is computationally lightweight. However, given the large number of patches in a single whole-slide image and the potentially large number of slides in a dataset, direct computation could still be prohibitive. Fortunately, computing persistence images is inherently embarrassingly parallelizable: each patch can be processed independently, and each slide can be analysed separately. This allows us to efficiently distribute computation across an HPC cluster, reducing processing time by a factor of proportional to the number of patches per slide, which could reach the thousands, particularly at high magnifications. In this paper, we leverage this parallelization, enabling us to generate bags of PIs and PCs for large-scale datasets in 2–3 hours.

In addition, since RGB patches and PIs are both 2D images, they can be processed using similar architectures--whether convolutional or attention-based--enabling a straightforward parallel stream for x'_i that mirrors the processing stream for x_i . For PCs, their analysis extends naturally from 2D images to 3D cubes, requiring only minor adjustments such as pre-trained weight mapping and modifying the initial layer to accommodate cross-channel convolution. This modularity ensures that the PMIL and cPMIL modules can be seamlessly integrated into any MIL approach with minimal code modifications.

Finally, by incorporating topological features, we introduce interpretable global representations that go far beyond pixel-level information. This improves the ability of pathologists to understand the underlying spatial structures of cellular components, such as nuclei and glandular formations, improving interpretability in clinical settings. This is in line with the growing demand for explainable AI in medical applications, which facilitates the adoption of deep learning tools in pathology.

3. Implementation details

For both the PMIL and cPMIL modules, the persistence information may represent H_0 or H_1 features, whether as bags of PIs or bags of PCs. Accordingly, we designate a dedicated

Table 1: Comparison of computational and memory overhead across a vanilla MIL model (CLAM), the proposed PMIL/cPMIL modules, and a hypothetical multi-scale RGB-based alternative. Values are reported per sample and normalized with respect to MIL. The proposed modules achieve substantial performance gains with only modest increases in FLOPs and memory, while avoiding the prohibitive costs associated with multi-resolution RGB processing. The symbol (\sim) indicates approximation.

Metric	MIL	PMIL/cPMIL	Multi-Scale RGB
Trainable Params	0.97M	1.3M (1.66 \times)	\sim 3.5M (4.4 \times)
FLOPs	1.59 MFLOPs	2.64 MFLOPs (1.66 \times)	\sim 6.5-8.0 MFLOPs (4-5 \times)
VRAM Usage	1 \times	\sim 1.6 \times	\sim 3.5 \times
Data RAM Usage	1 \times	\sim 1.2 \times (PMIL)/ \sim 2.5 \times (cPMIL)	\sim 6 \times

stream for each type, thus obtaining (c)PMIL-0 and (c)PMIL-1 streams. Each stream follows the same procedure described earlier, where the persistence image/cube is processed, and fused with the RGB features. The features are obtained through a ResNet50 model, pre-trained on ImageNet, without its last layer, to extract a feature vector of length d , where typically $d = 1024$. Herein, we deliberately exclude additional performance-enhancing techniques, such as data cleaning, stratified sampling, or pre-training the feature extractor on a histopathology-specific task. This ensures that the evaluation remains focused solely on the core design choices of our approach, preventing confounding effects from auxiliary optimizations.

For both modules, the final classification results from the fusion of the decision levels across the H_0 and H_1 streams. This fusion mechanism demonstrably improves classification performance by leveraging the complementary nature of topological features captured at different homology levels. To preserve the integrity of each stream, we keep their error gradients independent, allowing each to specialize in learning its respective homology without cross-interference. At this stage, the decision-level fusion integrates both streams into a cohesive and synergistic representation, as substantiated by our experimental findings.

4. Experiments

4.1 Experimental setup

We evaluated the two modules on the WSI classification task, comparing their performance against five widely used methods from the literature and two baseline approaches. The datasets used in this study are the SICAPv2^[31] dataset and Prostate cANcer graDe Assessment (PANDA-train) Challenge dataset.^[32] Among the two, SICAPv2 is considered a much smaller dataset, yet still representing a chance to evaluate our modules in the MIL setting.

SICAPv2 contains 155 WSIs, each assigned a Gleason Score of 0, 3, 4, or 5. The dataset includes approximately 18,000 labeled patches, but we discard the patch-level labels and use only the slide-level annotations. Unlike PANDA, SICAPv2 provides a single magnification level, making it applicable only to the PMIL module.

PANDA-train consists of 11,000 WSIs, divided into training, validation, and testing sets. Each slide is labeled on a Gleason Score scale from 0 to 5. The WSIs vary in size and include three magnification levels, starting at $0.5 \mu pp$. To ensure consistency with previous studies, RGB and PMIL experiments are conducted at $1.0 \mu pp$. Both PANDA and SICAPv2 are publicly available datasets: PANDA is provided via Kaggle under a CC BY-SA-NC 4.0 license, and SICAPv2 is released under a CC BY 4.0 license (non-commercial, attribution-only).

All training and testing were performed on completely separate sets, following a 10-fold cross-validation strategy. Training was conducted on a single NVIDIA A100 GPU.

In our experiments, we compare against several SOTA

methods, and two baseline methods, where the vanilla implementation of each operates in the RGB domain only. In detail, the MeanMIL and MaxMIL methods serve as simple baselines that aggregate patch-level features using mean and max pooling, respectively. We also evaluate against established attention-based models, including ABMIL,^[15] which uses a gated attention mechanism to assign importance scores to instances, and CLAM,^[17] which extends ABMIL by employing class-specific attention heads. Additionally, we include DSMIL,^[18] which identifies a critical instance to guide contrastive learning, ACMIL,^[20] which mitigates attention overfitting through top-K masking, and DTFD-MIL,^[19] which enhances multi-scale representation by forming pseudo-bags for hierarchical feature extraction. The PMIL module includes the bags of PIs, and the cPMIL module integrates the bags of PCs, as described in **Section 2. Methodology**. We report the results of slide-level classification using the accuracy (ACC) and area under the Receiver Operating Characteristic curve (AUC).

4.2 Results

The effect of implementing the topological features with magnification prior (cPMIL) or without (PMIL) are apparent by looking at the results of our experiments. Namely, **Fig. 5** compares between the different methods and the proposed modules. The positive effect of using the proposed modules is readily observed. Starting with the SICAPv2 dataset in **Fig. 5(a)**, the accuracy scores of five methods are improved upon the integration of the PMIL module out of seven, one observes no change, and one is slightly degraded. The AUC scores also undergo an improvement over four methods, whereas some degradation is observed over the last two methods. The levels of improvement/degradation are detailed in **Table 2**. Notably, even the simple MeanMIL baseline benefits significantly from the integration of PMIL, achieving a 4.9% increase in accuracy on SICAPv2-highlighting the utility of topological features even in the absence of complex attention mechanisms. Despite small degradation in some cases, it is evident that most methods benefit from the integration of topology information. A more consistent pattern can be inferred over the larger, more challenging PANDA dataset.

In the PANDA dataset, across all methods and over the two performance metrics, either one of the modules introduces an improvement over the vanilla implementation. The slight performance degradation incurred by incorporating the PMIL module into ABMIL is well compensated by a significant increase using the cPMIL method. The basic mean/max-pooling methods achieve significant performance improvement upon the integration of the modules.

More details can be gleaned from **Table 2**, where a significant improvement in the quality of cancer classification is observed. Specifically, the CLAM model,^[17] which appears to perform the best as shown in **Fig. 5(b)**, still achieves a 17% and a 7% increase in its accuracy and AUC, respectively, by augmenting it with the PMIL module. Similarly, the accuracy

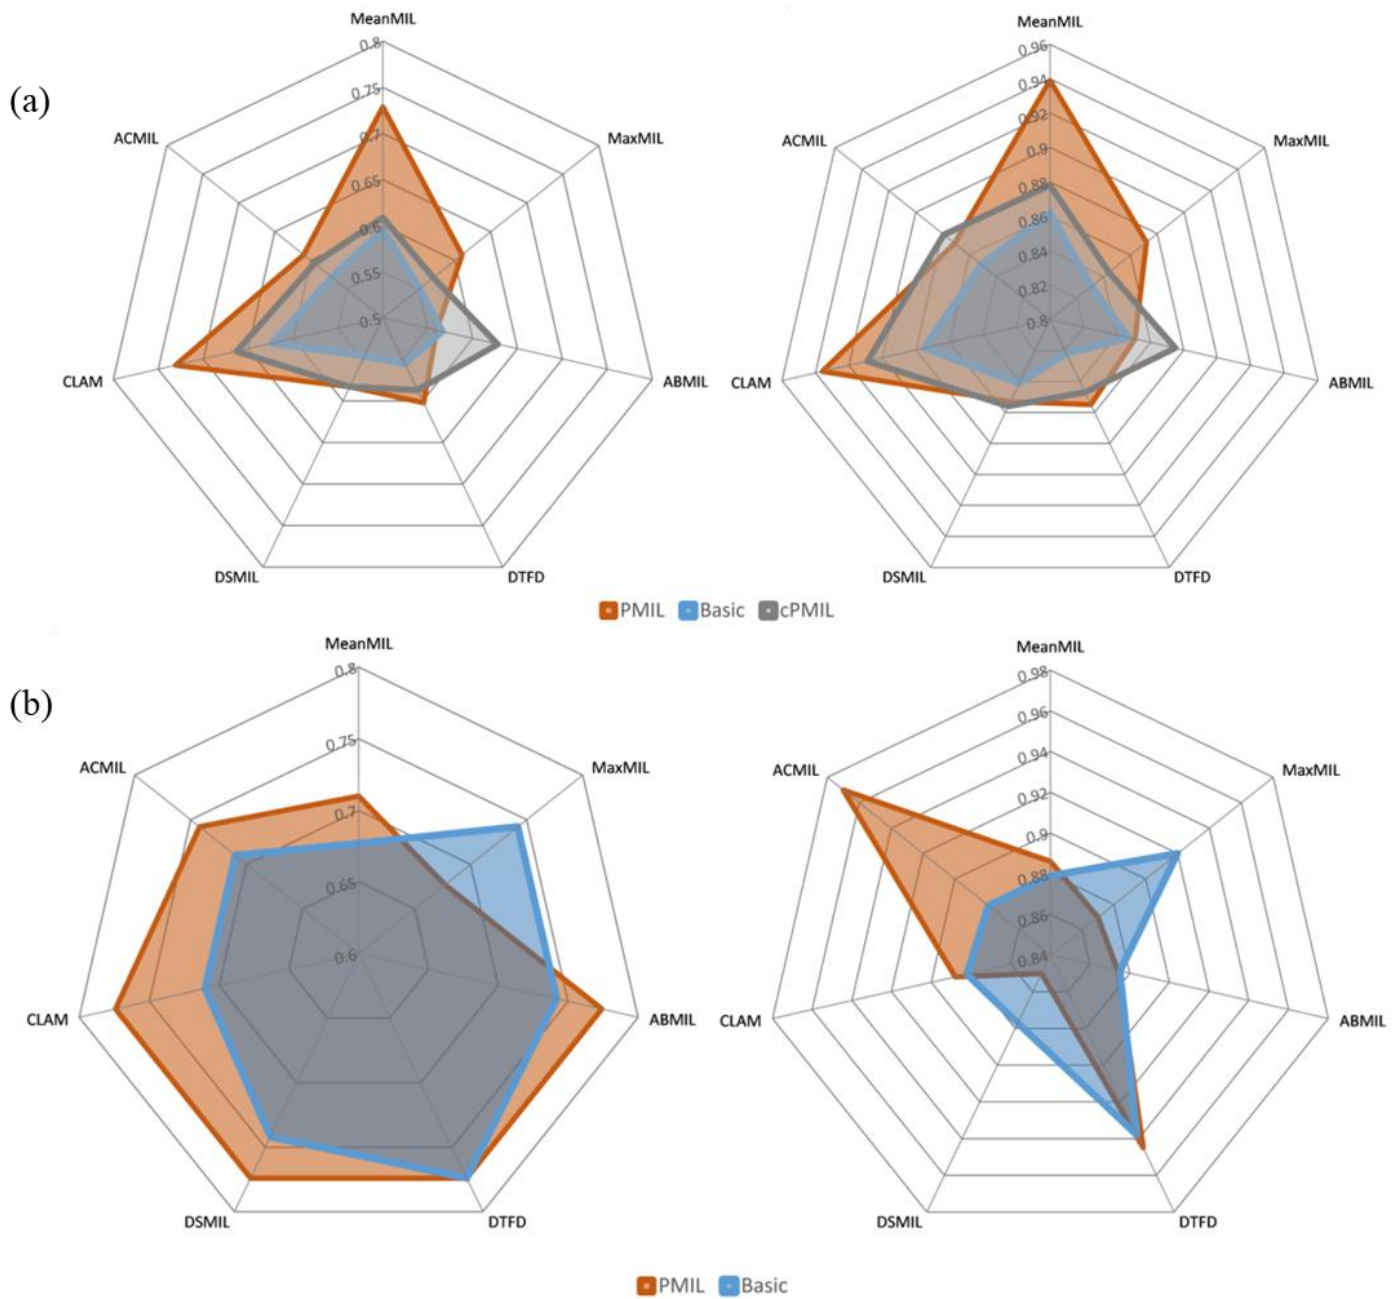


Fig. 5: The performance comparison between the vanilla implementation of several methods and the proposed modules over the ACC (left) and AUC (right) metrics. SICAPv2 is in (a) and PANDA is in (b).

Table 2: The percentage of performance change affected by the proposed modules over SOTA and baseline methods. Highest increase in each metric for each method is typed in Bold. Unavailable values are indicated by (-).

Method	PMIL		cPMIL	
	ACC	AUC	ACC	AUC
SICAPv2				
MeanMIL	↑4.9%	↑0.8%	-	-
MaxMIL	↓8.8%	↓5.4%	-	-
ABMIL	↑4.3%	↓0.1%	-	-
DTFD-MIL	0.0%	↑0.7%	-	-

Method	PMIL		cPMIL	
	ACC	AUC	ACC	AUC
DSMIL	↑4.3%	↓3.1%	-	-
CLAM	↑9.0%	↑0.7%	-	-
ACMIL	↑4.5%	↑10.4%	-	-
PANDA				
MeanMIL	↑22.3%	↑8.9%	↑2.2%	↑1.9%
MaxMIL	↑10.7%	↑4.9%	↑3.4%	↑1.6%
ABMIL	↓1.00%	↑0.5%	↑11%	↑3.3%
DTFD-MIL	↑8.50%	↑4.0%	↑5.8%	↑3.0%
DSMIL	↑6.80%	↑1.4%	↑6.8%	↑1.8%
CLAM	↑17.3%	↑7.0%	↑6.4%	↑3.9%
ACMIL	↑6.80%	↑2.4%	↑4.2%	↑3.3%

of the ABMIL^[15] model also improves by 10% and its AUC by 3% utilizing the cPMIL module. Both modules yield similar levels of improvement over the DSMIL^[18] method. Surprisingly, the second best method with a minute difference to CLAM^[17] (equipped with PMIL) is the basic MeanMIL, upon equipping it with PMIL. Despite its simplicity, this method outperforms the vanilla implementation of all SOTA methods.

A deeper insight is obtained by observing the characteristic of the patches attended to upon the integration of the topology features as can be seen in Fig. 6. In the first four columns, the pathologist assigned a Gleason grade 3, indicating moderate neoplastic differentiation of the glands. The PMIL module

appears to enhance the CLAM model's attention to the distinct topology of this grade, as evidenced by the warped, semi-circular glands that are highly attended to. In the last two columns, where the Gleason grades are 4 and 5, respectively, clusters of malignant cells forming disorganized, ill-defined glands receive more attention after implementing the PMIL module. Notably, this performance improvement, as evidenced by the experimental results, comes with no annotation overhead, *i.e.*, without requiring any additional samples in the training set, and a minimal increase in the computation and memory requirements, as discussed in Section 2.6 Optimizing performance, memory and computation.

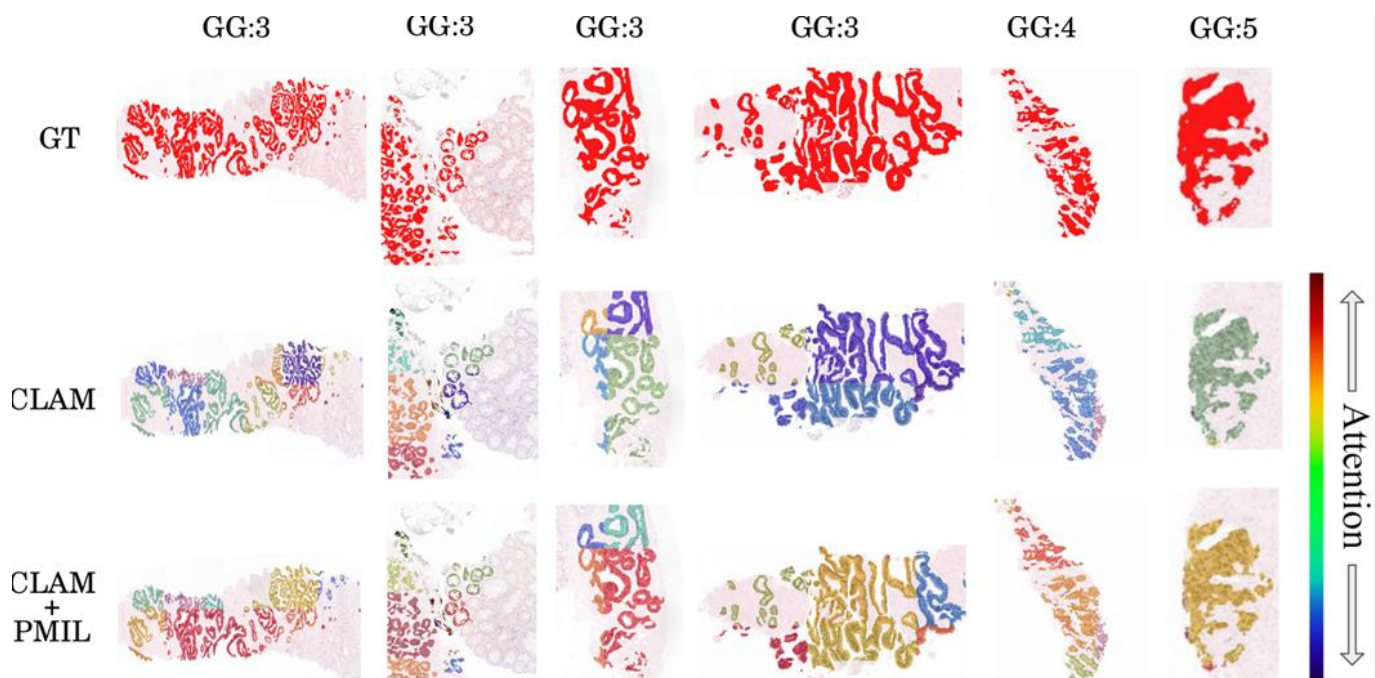


Fig. 6: Visual comparison of attention heatmaps at 1.0 μ pp. Primary patches informing the diagnosis of a pathologist (upper row) are compared with the highly attended patches of the vanilla CLAM Reproduced from^[17] method (middle row) and the highly attended patches after plugging the PMIL module (lower row).

5. Discussion

Despite increasing interest in integrating topological features into histopathology analysis, their potential remains largely untapped. This work aims to advance that frontier by tailoring TDA to the unique demands of computational pathology-specifically, through its integration with the MIL paradigm and adaptation to multi-magnification WSI data. In what follows, we elaborate on key strengths of the proposed modules, offer deeper insights into their behavior, and discuss limitations that may inform future work.

5.1 Quantifying nuclear density via average H_0

While H_1 features naturally capture glandular morphology, H_0 features provide a robust proxy for nuclear density and spatial organization. To empirically demonstrate this capability, Fig.




Fig. 7: Inverse of average H_0 persistence values plotted against representative tissue samples with varying nuclear clutter. A lower average H_0 (higher inverse) correlates with increased nuclear density, supporting its utility as a quantifiable measure of nuclear clutter.

5.2 Topological profiles across gleason grades

The interpretability of TDA is further illustrated by contrasting topological summaries across Gleason grades. As shown in Fig. 8, samples from Gleason Grade 3 display semi-structured glandular formations, resulting in distinctive, persistent H_1 features. In contrast, Grade 5 samples exhibit highly disrupted tissue architecture, minimal loop structures, and densely packed nuclei, leading to diminished H_1 signals and lower average H_0 persistence. These findings echo clinical expectations and align with prior work,^[9] reinforcing the diagnostic relevance of topological patterns. Importantly, such summaries enable global, interpretable characterizations of tissue morphology that go beyond pixel-level features and learned embeddings.

5.3 Effect of nuclei detection

The construction of persistent homology features in our framework begins with nuclei detection, as nuclei serve as critical biomarkers in the tumor microenvironment. Consequently, the accuracy of nuclei localization directly

visualizes the reciprocal average of H_0 persistence values across representative pathology samples. While the distinction between samples **iii**, **iv**, and **v** may be somewhat clear, the clutter levels of the first two samples are less describable visually, and may be prone to high subjectivity.

As such, the benefit of using topology is to project a quantifiable, objective description over the, otherwise, qualitative/visual attribute. Visually, it can be confirmed that a clear negative correlation exists between the average persistence in the H_0 profile and the density of nuclei, validating the use of average H_0 as a quantifiable biomarker for tissue density and organization. The benefit of this can be immense in segmentation models, conditioned diffusion synthesis, and other domains.

influences the quality of the resulting persistence diagrams. In this work, we deliberately opt for a simple thresholding-based method over more sophisticated, learning-based detectors such as HoVer-Net.^[25] This decision is motivated by computational efficiency: traditional thresholding techniques offer substantial speedups and are highly scalable, particularly when deployed on distributed or low-resource computing platforms.

To evaluate this trade-off, we compare the runtime performance across three histopathology datasets^[25,33-34] of varying sizes, summarized in Table 3. The results highlight the practicality of the simpler method in large-scale settings, where the processing burden of deep models may be prohibitive. Furthermore, topological representations such as persistence diagrams inherently exhibit robustness to mild detection noise. Their focus on persistent topological features allows the model to tolerate moderate false positives and negatives without significant degradation in downstream performance.

However, this design choice introduces limitations.




Fig. 8: Topological signatures of representative tissue samples from different Gleason Grade groups. Distinct patterns in H_0 and H_1 persistence reflect differences in nuclear density and glandular organization across cancer grades.

Table 3: Comparison of nuclei detection time across multiple histopathology datasets using HoVer-Net and a simple thresholding-based method.

Dataset	Tile count	Method	Time (s)
CoNSeP	41	HoVer-Net	130.8
		Simple	1.870
Lizard	238	HoVer-Net	561.0
		Simple	5.550
PanNuke	7901	HoVer-Net	2567
		Simple	0.588

Notably, simple thresholding lacks the capacity to differentiate between nuclear subtypes (e.g., lymphocytes, mitotic cells), which may hinder more granular topological analyses and limit biological interpretability. Although the current setup yielded strong results, incorporating more advanced detectors such as HoVer-Net in future iterations could enrich the extracted topological features with cell-type specificity. This could enable more refined modelling of histological structures, particularly in tasks involving immune infiltration or cell-specific spatial patterns.

5.4 Preliminary experimentation on homology levels

We conducted a series of controlled experiments to analyse the effect of using individual homology levels—PMIL-0 (H_0

only), PMIL-1 (H_1 only), and their feature-level fusion. For this part, we used the Camelyon16 dataset.^[35] Results, summarized in Table 4, highlight several key observations. First, feature-level fusion does not universally yield additive benefits; for instance, in the case of DTFD-MIL, the combined features underperform the better individual stream. Second, decision-level fusion consistently outperforms both feature-level fusion and individual streams. This suggests that H_0 and H_1 capture complementary, yet semantically distinct, information that is better preserved when error gradients are isolated during training.

Notably, feature-level fusion introduces a nontrivial memory overhead, particularly in the case of cPMIL where each cube contains multiple magnification slices. In contrast,

Table 4: Preliminary experimentation on the Camleyon16 dataset reveals the strength of decision-level fusion over individual streams and feature-level fusion.

Method	PMIL	PMIL-0	PMIL-1	Feature-Fusion
MeanMIL	0.729 0.939	0.716 0.939	0.711 0.931	-
MaxMIL	0.610 0.872	0.605 0.863	0.591 0.860	-
DTFD-MIL	0.602 0.855	0.576 0.845	0.586 0.845	0.582 0.841

our modular decision-level fusion allows each stream to operate independently, reducing memory usage and enabling more efficient optimization. This architectural decoupling not only promotes interpretability but also ensures scalability across increasingly complex WSI datasets.

5.5 Strengths, limitations and future directions

A major strength of the proposed PMIL and cPMIL modules lies in their modularity and seamless integration. These components require no architectural changes to existing MIL backbones and can be inserted with minimal code modifications. Importantly, they introduce topological descriptors—such as persistence of glandular loops or nuclear clustering—that are not only complementary to RGB features but also interpretable. This contributes to model trustworthiness, as these features can be visualized, verified by experts, and linked to known histopathological biomarkers. Such interpretability is particularly aligned with the growing emphasis on explainable AI in digital pathology, a prerequisite for clinical deployment.

In terms of comparative performance, our modules consistently improved accuracy and AUC when integrated with several state-of-the-art MIL methods, including CLAM, DSMIL, and ABMIL. Unlike these models, which rely solely on chromatic and morphological cues, PMIL and cPMIL offer robust, object-level modelling that captures spatial and structural irregularities overlooked by conventional attention or convolutional layers. This allows our approach to excel in noisy, artifact-laden, or poorly segmented regions—scenarios common in real-world pathology slides.

From a clinical perspective, the ability of topological features to reflect gland disruption and nuclear disorganization offers a quantifiable layer of decision support in tasks like Gleason grading. This could assist pathologists in making more consistent diagnoses, reduce inter-observer variability, and enable scalable computational screening—especially valuable in low-resource settings or high-throughput environments.

However, this study also has a few limitations. First, our experiments are limited to prostate cancer. While the modules were evaluated on glandular structures specific to prostate histopathology, the formulation of PMIL and cPMIL is agnostic to the tissue type and can be directly applied to other cancers where spatial topology is informative. In particular, cancers such as breast, colorectal, and lung adenocarcinomas, which also exhibit architectural complexity, are promising

candidates for extension. Future work should validate the generalizability of PMIL and cPMIL to other cancer types, such as breast and colorectal cancers, which also exhibit complex glandular architectures. Second, although we incorporated multi-magnification features, the filtration process relies on discrete magnification levels. More advanced techniques, such as continuous or differentiable filtration, could enhance both the expressiveness and end-to-end trainability of the modules. Lastly, the computation of persistence diagrams, while parallelizable and efficient for moderate datasets, may still pose scalability issues for ultra-large cohorts or resource-limited labs.

6. Conclusion

In this work, we have introduced PMIL and cPMIL, two versatile modules for modelling object-level information in histopathology images, and integrated them into five state-of-the-art MIL-based WSI classification pipelines and two baseline MIL methods. PMIL extracts topological features from pathology images by constructing Persistence Bags, which simulate the MIL framework's bag-of-instances structure. The cPMIL module extends this idea through the Bags of Persistence Cubes, to leverage the multi-magnification factor inherent in WSI datasets, mimicking the clinical practice of forming a final diagnosis based on differently magnified views. Through experiments on two challenging pathology datasets, we demonstrated that integrating topological features improves classification performance, as reflected in higher accuracy and AUC scores, and pronounced attention to glandular formations. This improvement is achieved without requiring additional annotated training data and with minimal memory and computational overhead. These results underscore the importance of topological features in histopathology analysis, an area still underexplored in computational pathology. To our knowledge, these modules are the first of their kind to apply TDA across entire WSIs.

As a future direction, our aim is to explore more efficient topological complexes, such as the Witness Complex,^[36] to improve computational efficiency. We hope this will help integrate topological analysis more seamlessly into future CPath research. Moreover, our current approach extracts a persistence image at each magnification level and concatenates the results into the persistence cube. We aim to explore a sparse and efficient alternative that explores the optimal filtration path in the (m, r) plane shown in Fig. 3.

Acknowledgments

This work is supported by research grants from the Advanced Technology Research Center Program (ASPIRE), Ref: AARE20-279, and the Terry Fox Foundation, Ref: I1037.

Conflict of Interest

There are no conflicts to declare.

Supporting Information

Not applicable.

CRedit Statement

Ahmad Obeid: Conceptualization, Data curation, Investigation, Methodology, Validation, Visualization, Writing – Original draft, Writing – Review & editing. **Sajid Javed:** Project administration, Writing – Review & editing. **Jorge Dias:** Project administration, Writing – Review & editing. **Ibrahim Abe M. Elfadel:** Conceptualization, Project administration, Writing – Original draft, Writing – Review & editing. **Naoufel Werghi:** Project administration, Funding acquisition, Resources, Validation, Writing – Review & editing.

References

- [1] D. G. Takale, A. V. Dhumane, P. N. Mahalle, T. Jadhav, P. P. Gawali, A. Buchade, Optimize deep learning model for intensive care of neurological disorders patients based on facial expression, *Engineered Science*, 2024, **32**, 1364, doi: 10.30919/es1364.
- [2] S. Ibrayev, B. Omarov, B. Amanov, Z. Momyunkulov, Development of a deep learning-enhanced lower-limb exoskeleton using electromyography data for post-neurovascular rehabilitation, *Engineered Science*, 2024, **31**, 1269, doi: 10.30919/es1269.
- [3] V. Kiruthik, S. Sathiya, R. M. M, K. Sakthidasan Sankaran, An intelligent machine learning approach for ovarian detection and classification system using ultrasonogram images, *Engineered Science*, 2023, **23**, 879, doi: 10.30919/es8d879.
- [4] A. Modi, B. Kishore, D. K. Shetty, V. P. Sharma, S. Ibrahim, R. Hunain, N. Usman, S. G. Nayak, S. Kumar, R. Paul, Role of artificial intelligence in detecting colonic polyps during intestinal endoscopy, *Engineered Science*, 2022, **20**, 25-33, doi: 10.30919/es8d697.
- [5] C. F. Loughrey, P. Fitzpatrick, N. Orr, A. Jurek-Loughrey, The topology of data: opportunities for cancer research, *Bioinformatics*, 2021, **37**, 3091-3098, doi: 10.1093/bioinformatics/btab553.
- [6] A. Bukkuri, N. Andor, I. K. Darcy, Applications of topological data analysis in oncology, *Frontiers in Artificial Intelligence*, 2021, **4**, 659037, doi: 10.3389/frai.2021.659037.
- [7] D. R. Chittajallu, N. Siekierski, S. Lee, S. Gerber, J. Beezley, D. Manthey, D. Gutman, L. Cooper, Vectorized persistent homology representations for characterizing glandular architecture in histology images, *IEEE 15th International Symposium on Biomedical Imaging (ISBI)*, Washington DC, USA, April 4-7, 2018, 232-235, doi: 10.1109/ISBI.2018.8363562.
- [8] T. Qaiser, Y.-W. Tsang, D. Taniyama, N. Sakamoto, K. Nakane, D. Epstein, N. Rajpoot, Fast and accurate tumor segmentation of histology images using persistent homology and deep convolutional features, *Medical Image Analysis*, 2019, **55**, 1-14, doi: 10.1016/j.media.2019.03.014.
- [9] P. Lawson, A. B. Sholl, J. Q. Brown, B. T. Fasy, C. Wenk, Persistent homology for the quantitative evaluation of architectural features in prostate cancer histology, *Scientific Reports*, 2019, **9**, 1139, doi: 10.1038/s41598-018-36798-y.
- [10] D. F. Gleason, Classification of prostatic carcinomas, *Cancer Chemotherapy Reports*, 1966, **50**, 125–128.
- [11] H. Wang, G. Huang, Z. Zhao, L. Cheng, A. Juncker-Jensen, M. L. Nagy, X. Lu, X. Zhang, D. Z. Chen, CCF-GNN: a unified model aggregating appearance, microenvironment, and topology for pathology image classification, *IEEE Transactions on Medical Imaging*, 2023, **42**, 3179-3193, doi: 10.1109/TMI.2023.3249343.
- [12] Y. Peng, H. Wang, M. Sonka, D. Z. Chen, PHG-net: persistent homology guided medical image classification, *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, Waikoloa, USA, January 3-8, 2024, 7568-7577, doi: 10.1109/WACV57701.2024.00741.
- [13] P. Bubenik, Statistical topological data analysis using persistence landscapes, *Journal of Machine Learning Research*, 2015, **16**, 77–102, doi: 10.5555/2789272.2789275.
- [14] H. Adams *et al.*, Persistence images: A stable vector representation of persistent homology, *Journal of Machine Learning Research*, 2017, **18**, 218-252, doi: 10.5555/3122009.3122017.
- [15] M. Ilse, J. Tomczak & M. Welling, Attention-based Deep Multiple Instance Learning, *Proceedings of the 35th International Conference on Machine Learning (ICML)*, 2018, 2127–2136, <https://proceedings.mlr.press/v80/ilse18a.html>.
- [16] P. O. Pinheiro, R. Collobert, From image-level to pixel-level labeling with Convolutional Networks, *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Boston, USA, June 7-12, 2015, 1713-1721, doi: 10.1109/CVPR.2015.7298780.
- [17] M. Y. Lu, D. F. K. Williamson, T. Y. Chen, R. J. Chen, M. Barbieri, F. Mahmood, Data-efficient and weakly supervised computational pathology on whole-slide images, *Nature Biomedical Engineering*, 2021, **5**, 555-570, doi: 10.1038/s41551-020-00682-w.
- [18] B. Li, Y. Li, K. W. Eliceiri, Dual-stream multiple instance learning network for whole slide image classification with self-supervised contrastive learning, *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Nashville, USA, June 20-25, 2021, 14313-14323, doi: 10.1109/cvpr46437.2021.01409.
- [19] H. Zhang, Y. Meng, Y. Zhao, Y. Qiao, X. Yang, S. E. Coupland, Y. Zheng, DTFD-MIL: double-tier feature distillation multiple instance learning for histopathology

- whole slide image classification, *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, New Orleans, USA, June 18-24, 2022, 18780-18790, doi: 10.1109/CVPR52688.2022.01824.
- [20] Y. Zhang, H. Li, Y. Sun, S. Zheng, C. Zhu, L. Yang, Attention-challenging multiple instance learning for whole slide image classification, *Computer Vision – ECCV 2024*, Springer Nature Switzerland, Cham, 2024, 125-143, doi: 10.1007/978-3-031-73668-1_8.
- [21] G. Carlsson, A. Zomorodian, The theory of multidimensional persistence, *Discrete & Computational Geometry*, 2009, **42**, 71-93, doi: 10.1007/s00454-009-9176-0.
- [22] O. Chapelle, V. Vapnik, O. Bousquet, S. Mukherjee, Choosing multiple parameters for support vector machines, *Machine Learning*, 2002, **46**, 131-159, doi: 10.1023/A:1012450327387.
- [23] O. Vipond, Multiparameter persistence landscapes, *Journal of Machine Learning Research*, 2020, **21**, 2262 - 2299, doi: 10.5555/3455716.3455777.
- [24] O. Vipond, J. A. Bull, P. S. Macklin, U. Tillmann, C. W. Pugh, H. M. Byrne, H. A. Harrington, Multiparameter persistent homology landscapes identify immune cell spatial patterns in tumors, *Proceedings of the National Academy of Sciences of the United States of America*, 2021, **118**, e2102166118, doi: 10.1073/pnas.2102166118.
- [25] S. Graham, Q. D. Vu, S. E. A. Raza, A. Azam, Y. W. Tsang, J. T. Kwak, N. Rajpoot, Hover-Net: Simultaneous segmentation and classification of nuclei in multi-tissue histology images, *Medical Image Analysis*, 2019, **58**, 101563, doi: 10.1016/j.media.2019.101563.
- [26] E. Reinhard, M. Adhikhmin, B. Gooch, P. Shirley, Color transfer between images, *IEEE Computer Graphics and Applications*, 2001, **21**, 34-41, doi: 10.1109/38.946629.
- [27] M. Macenko, M. Niethammer, J. S. Marron, D. Borland, J. T. Woosley, X. Guan, C. Schmitt, N. E. Thomas, A method for normalizing histology slides for quantitative analysis, 2009 *IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, Boston, USA, June 28 - July 1, 2009, 1107-1110, doi: 10.1109/ISBI.2009.5193250.
- [28] M. A. Kozminski, G. S. Palapattu, R. Mehra, J. S. Montgomery, A. Z. Weizer, T. A. Skolarus, B. K. Hollenbeck, D. C. Miller, C. He, S. Tomlins, J. E. Montie, F. Y. Feng, D. P. Wood, L. P. Kunju, T. M. Morgan, Understanding the relationship between tumor size, gland size, and disease aggressiveness in men with prostate cancer, *Urology*, 2014, **84**, 373-379, doi: 10.1016/j.urology.2014.03.037.
- [29] N. Y. Greenland, J. E. Cowan, B. A. Stohr, J. P. Simko, P. R. Carroll, E. Chan, Large cribriform glands (> 0.25 mm diameter) as a predictor of adverse pathology in men with Grade Group 2 prostate cancer, *Histopathology*, 2024, **84**, 614-623, doi: 10.1111/his.15102.
- [30] K. Nguyen, B. Sabata, A. K. Jain, Prostate cancer grading: Gland segmentation and structural features, *Pattern Recognition Letters*, 2012, **33**, 951-961, doi: 10.1016/j.patrec.2011.10.001.
- [31] J. Silva-Rodríguez, A. Colomer, M. A. Sales, R. Molina, V. Naranjo, Going deeper through the gleason scoring scale: an automatic end-to-end system for histology prostate grading and cribriform pattern detection, *Computer Methods and Programs in Biomedicine*, 2020, **195**, 105637, doi: 10.1016/j.cmpb.2020.105637.
- [32] W. Bulten, K. Kartasalo, P. C. Chen, P. Ström, H. Pinckaers, K. Nagpal, Y. Cai, D. F. Steiner, H. van Boven, R. Vink, C. Hulsbergen-van de Kaa, J. van der Laak, M. B. Amin, A. J. Evans, T. van der Kwast, R. Allan, P. A. Humphrey, H. Grönberg, H. Samaratunga, B. Delahunt, T. Tsuzuki, T. Häkkinen, L. Egevad, M. Demkin, S. Dane, F. Tan, M. Valkonen, G. S. Corrado, L. Peng, C. H. Mermel, P. Ruusuvoori, G. Litjens, M. Eklund, T. P. C. Consortium, A. Brilhante, A. Çakır, X. Farré, K. Geronatsiou, V. Molinié, G. Pereira, P. Roy, G. Saile, P. G. O. Salles, E. Schaafsma, J. Tschui, J. Billoch-Lima, E. M. Pereira, M. Zhou, S. He, S. Song, Q. Sun, H. Yoshihara, T. Yamaguchi, K. Ono, T. Shen, J. Ji, A. Roussel, K. Zhou, T. Chai, N. Weng, D. Grechka, M. V. Shugaev, R. Kiminya, V. Kovalev, D. Voynov, V. Malyshev, E. Lapo, M. Campos, N. Ota, S. Yamaoka, Y. Fujimoto, K. Yoshioka, J. Juvonen, M. Tukiainen, A. Karlsson, R. Guo, C.-L. Hsieh, I. Zubarev, H. S. T. Bukhar, W. Li, J. Li, W. Speier, C. Arnold, K. Kim, B. Bae, Y. W. Kim, H.-S. Lee, J. Park, Artificial intelligence for diagnosis and Gleason grading of prostate cancer: the PANDA challenge, *Nature Medicine*, 2022, **28**, 154-163, doi: 10.1038/s41591-021-01620-2.
- [33] S. Graham, M. Jahanifar, A. S. Azam, M. Nimir, Y. Tsang, K. C. Dodd, E. Hero, H. Sahota, A. Tank, K. Benes, N. Wahab, F. A. Minhas, S. E. Raza, H. Eldaly, H. K. Gopalakrishnan, D. R. Snead & N. M. Rajpoot, Lizard: A Large-Scale Dataset for Colonic Nuclear Instance Segmentation and Classification, *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, 2021, 684-693, doi: 10.1109/ICCVW54120.2021.00082.
- [34] J. Gamper, N. Alemi Koohbanani, K. Benet, A. Khuram, N. Rajpoot, PanNuke: an open pan-cancer histology dataset for nuclei instance segmentation and classification, *Digital Pathology*, Springer International Publishing, Cham, 2019, 11-19, doi: 10.1007/978-3-030-23937-4_2.
- [35] H. Yang, H.-Y. Zhou, C. Li, W. Huang, J. Liu, and S. Wang, *Generalizable vision-language pre-training for annotation-free pathology localization*, 2024, doi: 10.48550/arXiv.2401.02044.
- [36] E. Carlsson and J. Carlsson, *Alpha shapes in kernel density estimation*, 2023, doi: 10.48550/arXiv.2303.12213.

Publisher's Note: Engineered Science Publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access

This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International, which permits the use, sharing, adaptation, distribution and

reproduction in any medium or format, as long as appropriate credit to the original author(s) and the source is given by providing a link to the Creative Commons license. This usage for commercial purposes is not allowed. If modifications, adaptations or any other transformation were made, it is not allowed for distribution. The images or other third-party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <https://creativecommons.org/licenses/by-nc-nd/4.0/>.

©The Author(s) 2025.