



HAL
open science

A vision transformer-based framework for knowledge transfer from multi-modal to mono-modal lymphoma subtyping models

Bilel Guetarni, Feryal Windal, Halim Benhabiles, Marianne Petit, Romain Dubois, Emmanuelle Leteurtre, Dominique Collard

► To cite this version:

Bilel Guetarni, Feryal Windal, Halim Benhabiles, Marianne Petit, Romain Dubois, et al.. A vision transformer-based framework for knowledge transfer from multi-modal to mono-modal lymphoma subtyping models. 2024. hal-04456283

HAL Id: hal-04456283

<https://hal.science/hal-04456283>

Preprint submitted on 10 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

A vision transformer-based framework for knowledge transfer from multi-modal to mono-modal lymphoma subtyping models

Bilel Guetarni¹, Feryal Windal¹, Halim Benhabiles², Marianne Petit³, Romain Dubois³,
Emmanuelle Leteurtre³, and Dominique Collard⁴

¹Junia, UMR 8520, University of Lille, CNRS, Centrale Lille, University of Polytechnique
Hauts-de-France, F-59000, Lille, France

²IMT Nord Europe, Institut Mines-Télécom, Univ. Lille, Centre for Digital Systems, F-59000
Lille, France

³Department of Pathology of the University Hospital of Lille, France

⁴LIMMS/CNRS-IIS, IRL 2820, The University of Tokyo, and is the E-Health Chairholder of
Yncrea Méditerranée, Yncrea Ouest, Junia, Toulon, Brest, Lille, France

Abstract

Determining lymphoma subtypes is a crucial step for better patient treatment targeting to potentially increase their survival chances. In this context, the existing gold standard diagnosis method, which relies on gene expression technology, is highly expensive and time-consuming, making it less accessible. Although alternative diagnosis methods based on IHC (immunohistochemistry) technologies exist (recommended by the WHO), they still suffer from similar limitations and are less accurate. Whole Slide Image (WSI) analysis using deep learning models has shown promising potential for cancer diagnosis, that could offer cost-effective and faster alternatives to existing methods. In this work, we propose a vision transformer-based framework for distinguishing DLBCL (Diffuse Large B-Cell Lymphoma) cancer subtypes from high-resolution WSIs. To this end, we introduce a multi-modal architecture to train a classifier model from various WSI modalities. We then leverage this model through a knowledge distillation process to efficiently guide the learning of a mono-modal classifier. Our experimental study conducted on a lymphoma dataset of 157 patients shows the promising performance of our mono-modal classification model, outperforming six recent state-of-the-art methods. In addition, the power-law curve, estimated on our experimental data, suggests that with more training data

from a reasonable number of additional patients, our model could achieve competitive diagnosis accuracy with IHC technologies. Furthermore, the efficiency of our framework is confirmed through an additional experimental study on an external breast cancer dataset (BCI dataset).

1 Introduction

Diffuse Large B-Cell Lymphoma (DLBCL) is a cancer of the immune system cells with a significant worldwide impact. More specifically, in 2020, 544,352 new cases have been diagnosed and 259,793 deaths have been recorded [1]. Additionally, projections indicate a 7% increase in both new cases and deaths in Western Europe from 2020 to 2025 [2]. It has been shown that distinguishing between DLBCL subtypes, namely ABC (activated B-cell-like) and GCB (germinal-center B-cell-like), is crucial for treatment selection and the vital prognosis of the patient [3]. Indeed, clinical trials using new targeted treatments have shown promising results in terms of efficiency and survival chance improvement, notably for the patients with the most aggressive subtype (ABC) [4, 5].

Currently, the gold standard method for DLBCL subtyping is a molecular biology method named RT-MLPA (*reverse transcriptase multiplex ligation-dependent probe*

amplification [3]). Nevertheless, its high cost (600€¹) and longtime response (35 days¹) have maintained the exploitation of other diagnostic methods, notably the microscopic examination of the patient tissue by a pathologist.

The tissue examination can be performed on two types of markers, namely IHC² (immunohistochemistry) and HES (hematoxylin-eosin-safran) markers. In the case of IHC, the expression levels of different proteins are quantified by the pathologist, which permits the determination of the DLBCL subtype based on the Hans algorithm [6]. Although the response time of this latter method is relatively fast (2 to 3 days¹) and its cost (90€¹) is reasonable compared to the gold standard, it has an error rate of 20% (misclassified cases). Additionally, it requires the involvement of the pathologist for the tissue examination which represents a fastidious task, notably in daily routines. In the case of HES markers, the pathologists analyze the morphology and structural organization of the tissue cells. However, due to the high similarity of these features between the two subtypes and the lack of objective criteria, pathologists struggle to reach an agreement on their decision. Compared to the IHC, the HES analysis is less expensive (it does not require protein markers) but suffers from a higher error rate (30-35%) with respect to the gold standard.

With the rapid increase of data in medical centers, there has been a growing interest in applying deep learning methods to medical image analysis [7, 8, 9]. In the field of histopathology, the slide of a tissue sample can be digitized into high-resolution images named Whole Slide Images (WSI). Deep learning methods have shown unprecedented performance on many automatic tasks on these images [10]: detection of several cancers (e.g., breast, lung, colon, etc. [11, 12, 13]), nuclei segmentation [14] or case/image similarity [15, 16, 17]. Although several works have addressed lymphoma subtyping [18, 19, 20, 21, 22, 23], to the best of our knowledge, there is no prior work on automatic differentiation of DLBCL molecular subtypes based on WSI analysis. Among the main reasons that could explain this observation is the scarcity of publicly available annotated datasets (data protection and privacy). In addition, the characterization of the two subtypes of DLBCL from WSI is a challenging task. Indeed, the diversity of the cancer location in the patient’s body and the complex nature of the tissue lead to high intra-class variability for each subtype. Besides, the extremely high resolution of WSI (100,000 × 100,000 pixels on average) prevents

the direct application of existing pretrained CNN architectures and involves highly demanding computational requirements.

In this work, we propose a deep learning-based framework for distinguishing between the ABC and GCB subtypes of DLBCL cancer from high-resolution WSI. Our framework involves the development of a multi-modal deep architecture, serving as a teacher model for classifying the cancer subtypes using various WSI modalities, namely IHC-BCL6, IHC-CD10, IHC-MUM1 and HES. Additionally, we leverage the concept of knowledge distillation to build a mono-modal classifier, referred to as the student model, which shows promising performance for differentiating between the subtypes using a single modality, namely HES. It is worth mentioning that having a subtyping method based only on HES modality is of high interest for the health care system (pathologists, patients, laboratories). Indeed, removing the three IHC modalities will permit to (i) lighten the data acquisition process in the laboratories, (ii) accelerate patient diagnosis, and (iii) avoid the costs related to the production of these modalities.

The main contributions of this work are the following:

- We exploit the concept of knowledge distillation for efficiently driving the learning of a classifier following a *4-to-1 WSI modality* training process. In this sense, the teacher model is used to support the student one to extract the most relevant features for DLBCL subtype characterization. In addition, we propose a generic design of the framework, enabling its exploitation for any classification task that requires knowledge transfer from a multi-modal classifier to a mono-modal one.
- We designed and developed a deep architecture tailored for handling multi-modal inputs (four modalities of WSI) in a one-shot pass. To deal with the extremely high resolution (100,000 × 100,000 px) of these inputs, we incorporated in our architecture a ViT encoder that permits processing them in the form of a sequence of patches.
- We designed and developed a module offering a multi-modal features fusion mechanism to efficiently learn a discriminative representation of the two cancer subtypes.
- We conducted a comprehensive experimental study to demonstrate the efficiency of our method and its superiority compared to state-of-the-art approaches.

¹estimation from the Lille University Hospital

²Alternative diagnosis method recommended by the WHO when RT-MLPA is not available.

2 Related works

Recently, Guidolin et al. [24] have proposed an experimental study on 30 DLBCL patients, comprising 15 ABC and 15 GCB cases. They aimed to quantify the spatial distribution of cells in WSIs using statistical measures, using five IHC protein markers for staining. As raised by the authors, their experimental study has permitted to highlight that the combination of spatial statistics-derived parameters can be exploited for distinguishing between ABC and GCB. However, the authors did not propose any automatic subtyping method.

Although the research on automated DLBCL subtyping is scarce, there is an extensive literature on utilizing deep neural networks for cancer classification from WSI [10, 25, 26, 11, 13, 27, 28, 29]. Among these works, several ones have addressed the differentiation between DLBCL and other lymphoma types [18, 19, 20, 21, 22, 23]. For example, Hashimoto et al. [22] proposed a CNN-based classification method to distinguish between three lymphoma cancers, namely DLBCL, AITL and CHL, from HES WSIs. To deal with the high resolution of these images, they made the choice to subdivide them into a set of patches and adopted a Multiple Instance Learning (MIL) approach to train the final classifier. To improve the performance of the classifier, they proposed a methodology that allows for the selection of the most relevant WSIs from the dataset during the training stage. More specifically, they defined a typicality criterion that allows for the measurement of the relevance of each WSI in terms of cancer class representativeness.

In addition to the task of cancer classification from WSI, other works proposed a higher level of cancer analysis-related applications, namely patient survival prediction [30, 28]. In the case of [30], the authors designed a multi-modal graph neural network architecture to train a model to predict survival chances in cases of bladder and ovarian cancer from WSI IHC modalities. As raised by the authors, using a GNN enables the extraction of cell-centric features and processing of the WSI in a single forward pass. In the case of Chen et al. [28], they proposed a backbone composed of a series of vision transformer units to encode multi-scale region features. To this end, they choose to feed these units by exploiting the levels of a pyramidal representation of the WSI, where each level corresponds to a subdivision into patches of customized size. To improve the performance of their model, they pretrained it by adopting a self-supervised learning strategy.

Other works have been proposed to address the diagnosis of different diseases from WSIs such as colorectal polyps classification [31], neurodegenerative disease diagnosis [32] or virus infected cell identification [33]. For

instance, the method proposed by Dwivedi et al. [34] performs liver disease severity scoring from two WSIs modalities, namely Masson’s trichrome and HES. For each modality, a CNN generates a heatmap of the WSI which is subsequently transformed into a graph representation. Each heatmap is then processed by a graph neural network (GNN) to produce a vector that captures the characteristics of the WSI. The disease severity score is predicted from the concatenation of the two modality vectors.

While previous studies [35, 36, 37, 38] have applied knowledge distillation for multi-modal to mono-modal knowledge transfer in the context of medical image segmentation, our work stands out as we have developed a unique approach tailored for tissue classification, a distinct task from segmentation, which handles high-resolution images.

3 Methods

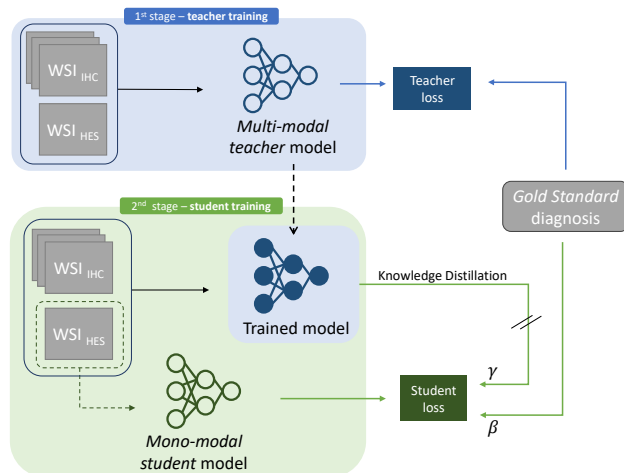


Figure 1: Knowledge Distillation for multi-to-mono modal WSI subtyping model.

As illustrated in Fig. 1, our framework designed for cancer subtyping from WSIs includes two classification models that are built by following a two-stage training process. In the first stage, a multi-modal teacher classifier is trained on HES and IHC modalities. In the second stage, the teacher acts through a knowledge distillation process as a robust supervision source for training a mono-modal student classifier on the HES modality.

Our deep learning architecture is summarized in Fig. 2, and we next detail each component.

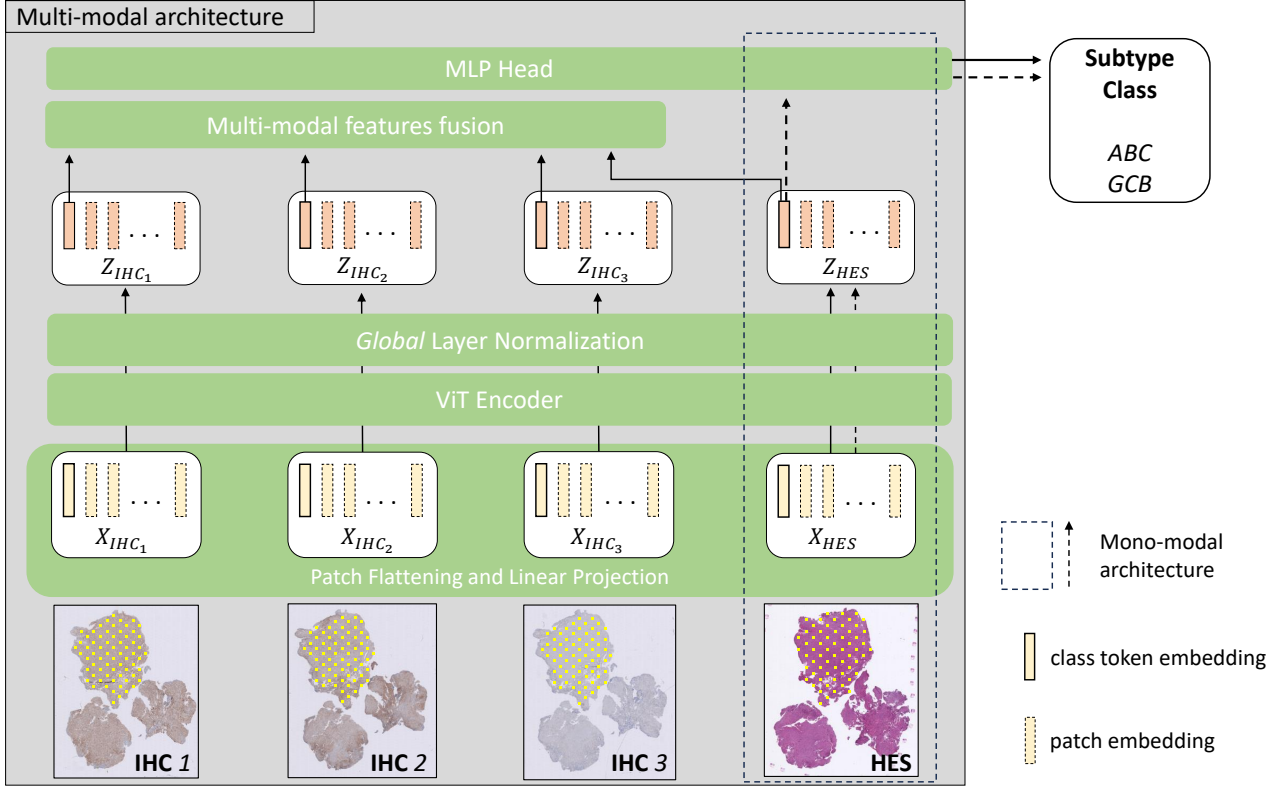


Figure 2: Our multi-modal architecture (teacher) takes as input a bag of sequences uniformly sampled from the same region across all WSI modalities (yellow squares). Its multi-modal features fusion mechanism produces a bag representation for multi-modal WSIs classification. Our mono-modal architecture is presented within the dotted bounding-box on the right.

3.1 From multi-slides input to single prediction

In this section, we delve into the specifics of how our methodology processes WSIs stained with HES as well as IHC.

To avoid excessively complex computations related to the size of the regions, we adopted a region-based sampling approach. It involves extracting matching regions of interest between all WSI modalities (see Fig. 2). For each region R and each modality m , we extract from the region, RGB patches $x \in \mathbb{R}^{p \times p \times 3}$, with size $p \times p$, and we constructed sequences (X) of these patches as follows:

Let N denote the number of patches extracted from R , and S represents the number of patches in a sequence such that $X \in \mathbb{R}^{S \times p \times p \times 3}$. We systematically extract sequences by randomly sampling S patches uniformly distributed in R , without replacement. This process is repeated until $N < S$ (i.e., there are less than S patches left in the region). Let R_m be the set of all sequences extracted from region R of modality m . Let B be a bag formed by sequences from all modalities, represented

as:

$$B := \{X_m, \forall m \in M\} \quad (1)$$

where $X_m \in R_m$ represents a sequence from modality m , and M is the set of modalities. Note that the sequences are sampled without replacement from each modality. Indeed, when all sequences of one modality have been used, no more bags can be formed. The bag B serves as the input to the multi-modal model (i.e., teacher model).

Given that regions are represented as series of patches, we utilize the ViT architecture [39] as our feature extractor. Furthermore, this architecture has the capacity to learn inter-element relationships via its self-attention mechanism, which is of high interest in the case of WSI patches [28, 27]. The ability to understand and leverage these relationships is crucial to learn global representation of WSIs.

To classify an image, the ViT architecture [39] transforms the image into a sequence of patches. These patches are subsequently flattened and projected by a linear layer into vectors known as patch embeddings. In addition to this sequence, a learnable class token

is prepended, serving as the image embedding post-encoder for the purpose of classification. Our approach adheres to a similar structure: we flatten the patches of the sequences, project them with a linear layer, and prepend a class token. However, in our scenario, we are working with a bag of sequences representing the same object but manifesting across different modalities (i.e., stains). To account for variations in color distribution among these modalities, we employ a separate projection layer for each one. The learnable class token embedding, however, is shared across all modalities. The bag sequences are independently encoded by a shared ViT encoder and subsequently normalized with a Global Layer Normalization (GLN) [40]. All sequences X_m of the bag B are encoded and normalized simultaneously:

$$Z_m = GLN(Encoder(X_m)) \quad (2)$$

where $Encoder$ is the ViT encoder (including patch flattening and linear projection) and GLN is the Global Layer Normalization. We use the class token embeddings from every sequence after this GLN as input to our multi-modal features fusion mechanism, which produces the bag representation. In our experimental study, we found that adding a softmax activation after the encoder’s MLP improves the validation results. In what follows, we refer to this modified version of the encoder unless otherwise stated.

3.2 Multi-modal features fusion

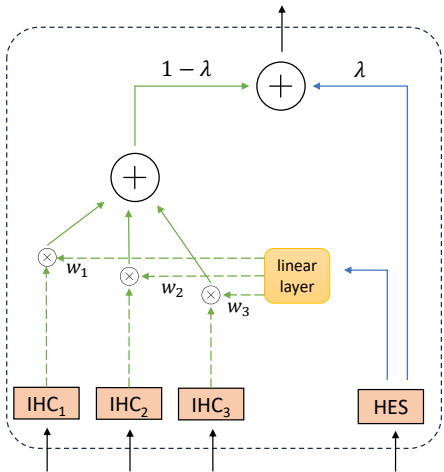


Figure 3: Our multi-modal features fusion mechanism.

We design our multi-modal features fusion mechanism (Fig. 3) to fuse information from IHC and HES

embeddings. Because every modality is not equally important for different samples, we use a learnable attention weights approach to fuse IHC features before fusing it with HES. For this, we consider HES as a query that is used to select information from IHC embeddings. Whereas the current popular attention mechanism is defined as the dot-product attention between a query and keys [39], in our case, different modalities are involved, making the dot-product-based attention unsuitable. When the query and keys come from different space regions, it becomes less relevant to use the dot-product between the representation of the query and keys as a similarity indicator. In our case, this is particularly evident as HES images exhibit a significantly different color distribution than IHC images (see inputs in Fig. 2). In this regard, the attention weights are not calculated as a dot-product between the query and the keys but are predicted by a linear layer from the query.

Let z_{HES}^0 be the HES class token embedding (initialized with the shared learnable class token) following the GLN. Similarly, let $z_{IHC_i}^0$ be the class token embedding from IHC i . The multi-modal features fusion starts by predicting IHC attention weights with a linear layer from the HES class token embedding:

$$w = Az_{HES}^0 + b \quad (3)$$

where A and b are the weights and bias of the linear layer respectively. This linear layer outputs a k -dimensional vector, with k the number of IHC stains ($k = 3$ in our case). We then perform a weighted sum of the IHC class token embeddings using these weights and add it to the HES one:

$$z_B = \lambda z_{HES}^0 + (1 - \lambda) \sum_{i=1}^k w_i z_{IHC_i}^0 \quad (4)$$

where λ is a trade-off coefficient between HES and IHC features (which is manually set in our experiments). The output of multi-modal features fusion is a vector representation of the bag B . This vector is then fed to a classification head for subtype classification.

In the first stage, the multi-modal model is trained to minimize the cross-entropy loss with the ground-truth of the WSI from which the bag is sampled.

3.3 Knowledge Distillation from multi to mono modality WSI classification

We utilize KD to transform the multi-modal model, initially trained with both IHC and HES data (teacher), into a mono-modal one that is solely reliant on HES data (student). During the training of the student,

we feed the teacher the bag of sequences, while the student predicts the subtype from the HES sequence only. Because the student doesn't make use of IHC we remove the multi-modal features fusion mechanism; indeed, the HES class token (z_{HES}^0) after the LN is then directly fed to the classification head. After the teacher model has finished training, it produces outputs that are the corresponding probability predictions for each class. Subsequently, the student model begins its training stage. The student model's progress is supported by leveraging the outputs from the teacher model in addition to the ground-truth. The objective during the training phase of the student model is to minimize the cross-entropy between its own predictions and those of the teacher model, as well as between its predictions and the actual target. In this study, we implement knowledge distillation in the form of hard distillation, indicating that we treat the teacher model's prediction as a hard target for the student model's training. If f^t and f^s denote the teacher and student models respectively, X^t and X^s their corresponding inputs and y the ground-truth, the student model loss function is defined as:

$$Loss = \beta \mathcal{L}(y, f_s(X^s)) + \gamma \mathcal{L}(f_t(X^t), f_s(X^s)) \quad (5)$$

where \mathcal{L} is the cross-entropy loss function, β and γ are coefficients for the ground-truth and teacher supervisions respectively. We convert the teacher's output to hard labels and apply label smoothing [41] with a parameter of $\alpha = 0.1$ to both the teacher's prediction and the true label.

4 Experiments

4.1 Data description and processing

We conduct our experiments on a retrospectively collected cohort of WSIs from patients of the Lille University Hospital diagnosed with DLBCL between 2010 and 2020. The collection has been meticulously inspected by a pair of pathologists, senior and junior. Subtyping of the patients was conducted in three distinct ways:

- pathologist diagnosis based on visual examination of the HES WSI
- application of the Hans algorithm [6]
- execution of RT-MLPA [42]

However, given that RT-MLPA is recognized as the gold standard for DLBCL subtyping, we take its result as the ground-truth for the subtyping of patients.

Our initial dataset consisted of 204 patients, including 94 diagnosed with the ABC subtype, 66 with GCB, and 44 who did not fall into either category (being either *unclassifiable* or *primary mediastinal large B-cell lymphoma*). We have excluded these 44 patients from our dataset. Additionally, 3 patients were removed due to the absence of discernible cancerous tissue in their HES WSI, which is crucial for accurate analysis. For the training phase, we selected 93 patients out of the remaining 157 subjects. For these patients, we were able to successfully match the HES WSI with their corresponding IHC WSIs (BCL6, CD10, and MUM1). During the first stage, the teacher is trained with the entire bag of sequences containing all modalities, while in the second stage (KD), the student is trained with the HES sequences only. For the testing phase, we include the remaining portion of the dataset, which consists of 64 patients with a single HES WSI. The train set patients are randomly separated into training and validation according to a 80:20 split.

All slides were scanned with a Ventana iScan HT at $\times 40$ magnification ($0.25\mu\text{m}$ per pixel). For data pre-processing, the tissue was separated from the background through gaussian blurring and Otsu thresholding segmentation. For each bag and each modality, we consider sequences of patches with $S = 256$ and $p = 32$ extracted at $\times 20$ magnification. Our dataset covers 38 different tissue types such as: lymph node (42%), brain (13%), mediastinum (5%), soft tissue (5%), bone (4%), etc. This large diversity of tissue location results in very heterogeneous image characteristics among DLBCL images.

Table 1 provides a summary of the distribution of sequences by modality and the distribution of patients in the training and testing splits.

Table 1: Data summary: number of patients for each set. Training patients were split into 80:20 for train and validation.

	train		test		total
	ABC	GCB	ABC	GCB	
sequences					
<i>BCL6</i>	83K	33K			116K
<i>CD10</i>	82K	33K			115K
<i>MUM1</i>	80K	35K			115K
<i>HES</i>	78K	35K	34K	32K	179K
patients	59	34	33	31	157

4.2 Experimental settings

For both teacher and student models we use a ViT encoder ($d=128$, $n=3$, $h=2$). The classification head is a trainable linear layer with two output neurons (one for each subtype), and a softmax activation to transform its output to a probability distribution over the subtypes. Both the teacher and student models are trained from scratch. To determine a patient subtype, we apply major voting on all the sequences prediction coming from this patient. More specifically, we exploit our sequence-based classification model to predict a class for each sequence of the WSI (each sequence having a length of 256 patches), and then the subtype corresponds to the most dominant class over all the predictions. The teacher is trained with a batch of 32, where each sample of the batch is a bag of sequences (B). We use an initial learning rate of 8.10^{-5} with a learning rate decay of 0.5 every 5 epochs. The student is trained with a batch of 64 and a constant learning rate of 10^{-5} . The loss coefficients of the student were set according to best validation accuracy, with $\beta = 0.5$ and $\gamma = 0.5$. Similarly, the trade-off coefficient of the multi-modal features fusion of the teacher λ was set to 0.5. For both models, we adopt the Adam optimizer with clipping gradient (max norm of 3 for teacher and 5 for students), and apply label smoothing ($\alpha = 0.1$) to the ground-truth. The models are trained for 100 epochs, and the weights with the lowest validation loss are saved. We apply class balancing by trimming the dominant class, it has shown to improve the validation results. All models are trained and tested with PyTorch 1.8 on 4 NVIDIA Tesla V100 GPUs, and the metrics are calculated with the scikit-learn (v. 1.1.2) package. For 100 epochs, the teacher and student training takes 17 and 16 hours respectively.

To analyze and compare the performance of our method we use the precision (PRE), recall (REC) and total accuracy (ACC) which corresponds to the patients correctly classified.

4.3 Comparison with state-of-the-art methods and models

We conduct a comparative analysis with state-of-the-art methods in WSI classification. For this, we have selected several notable works for this purpose including:

- HIPT [28] is composed of hierarchical ViTs with self-supervised pretraining on 10,678 WSI from the TCGA open dataset.
- RSP-CR [43] pretrained a ResNet-18 with self-supervision on 69 WSIs. The fine-tuning phase

involves consistency regularization.

- CTransPath [45] uses a hybrid CNN-Transformer architecture with self-supervised pretraining on 32,220 WSIs from the TCGA and PAIP datasets. The technique employs semantically relevant contrastive learning for unsupervised pretraining.
- MS-DA-MIL [44] introduces a multi-scale domain adversarial training approach for WSI classification using a CNN with MIL.
- DTFD-MIL [13] extends the MIL approach for WSI classification by introducing pseudo-bags with attention-based MIL.
- KimiaNet [17] fine-tuned an ImageNet pretrained DenseNet-121 on 7,126 WSIs from the TCGA dataset with 32 cancer classes.

For training, we use the pretrained weights provided by the authors for HIPT, RSP-CR and CTransPath, as well as the fine-tuned weights for KimiaNet. For the later, we reproduce the same downstream approach in [17] (i.e. we fit a cubic SVM on top of its features). We apply the specific hyperparameters of all methods from their respective paper. We have evaluated the six methods on our dataset, ensuring the same data split rules as our method. Results are presented in Table 2.

According to the metrics we have evaluated, our approach demonstrates superior or comparable results in most cases, when measured against the state-of-the-art methodologies. While CTransPath displays a marginally better ability to correctly identify GCB patients, with a 3% higher recall rate, our model consistently delivers a more balanced performance. Our F-score, a measure that includes both precision (how many selected items are relevant) and recall (how many relevant items are selected), is higher by 2%. This implies that our method is better at both correctly identifying positive cases and minimizing the number of false positives, providing a more reliable and accurate overall performance. Despite their pretraining on large datasets, models like CTransPath, HIPT, or KimiaNet are still outperformed by our method in effectively distinguishing between the two subtypes, with respective accuracy gaps of +5%, +23%, and +14%. DTFD-MIL, which shows the best results among all comparison methods, is outperformed by our method by +3%. Furthermore, it requires on average $25\times$ more floating operations and contain $10\times$ more parameters than our student model. We observe that most methods tend to misclassify GCB samples. Indeed, among the six compared methods, four ones (namely HIPT, RSP-CR, MS-DA-MIL and KimiaNet) gave for the GCB class a

Table 2: Comparison of the Student model with state-of-the-art methods for WSI classification on test set (64 patients). We estimated the number of floating operations (number of op.) for a WSI diagnosis based on 5 runs.

	ABC		GCB		ACC \uparrow	number of op. ($\times 10^{12}$)	params.
	PRE	REC	PREC	REC			
HIPT* [28]	0.52	1.00	0.00	0.00	0.52	–	25M
RSP+CR* [43]	0.54	0.94	0.71	0.16	0.56	30.8	11.8M
MS-DA-MI [44]	0.58	0.76	0.62	0.42	0.59	39	67M
KimiaNet* [17]	0.58	0.85	0.69	0.35	0.61	11.5	7M
CTransPath* [45]	0.71	0.73	0.70	0.68	0.70	25	28M
DTFD-MIL [13]	0.68	0.85	0.78	0.58	0.72	18.6	9M
ours (student)	0.72	0.85	0.80	0.65	0.75	0.5	0.89M

* indicates pretrained models

the abbreviations are: precision (PRE), recall (REC) and accuracy (ACC)

recall value lower than 0.5. KimiaNet and DTFD-MIL display a comparable ABC recall than our method (0.85 each), but have inferior GCB recall (0.35, 0.58 and 0.65 respectively). Additionally, we observe that our model is very light compared to other methods, with less than one million parameters, whereas other models contain between 7 and 67 million parameters.

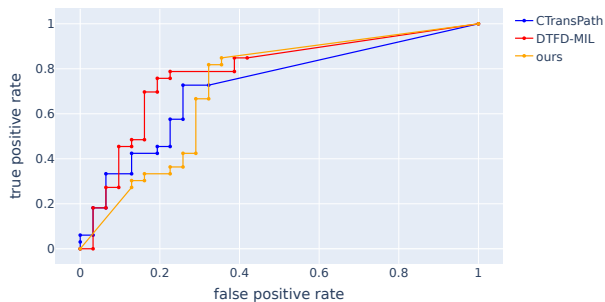
These models face challenges in accurately distinguishing between DLBCL subtypes due to the specific nature of this task. Indeed, because the task is highly complex due to the presence of visually similar features among the subtypes, these methods also do not escape this complexity and struggle to effectively capture or adapt to it during their training process. On the other hand, by leveraging different data modalities, our teacher model is able to learn a richer set of features better suited to the task at hand. The teacher model then provides a crucial supervision to guide the student model to discover the most relevant features for the given task.

Furthermore, we generate ROC curves, for each class separately, for our student model and the two methods in Table 2 with an accuracy higher than 70%. The curves for the ABC class are shown in Figure 4a. We observe that DTFD-MIL shows relatively higher true positive rates (TPR) for low false positive rates (FPR); see FPR in the range [0, 0.3] in the Figure 4a. However, our method shows the best TPR when the FPR exceeds 0.3. To more explicitly illustrates the sensitivity of the methods with respect to the classification threshold, we have calculated the accuracy of the classifiers for this class, across several thresholds (in the range [0.5, 1]), displayed in Figure 4b. We observe that at a threshold of 0.5, our method shows the best accuracy (0.75), however, contrary to DTFD-MIL, its accuracy decreases with higher thresholds. Overall, our method still ex-

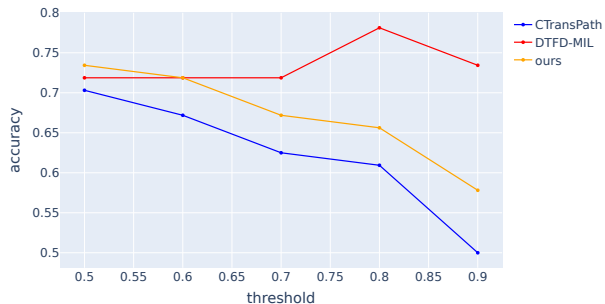
hibits for this class the best F1-score (0.78) compared to DTFD-MIL (0.76) and CTransPath (0.72). Similarly, the ROC curves for the GCB class are shown in Figure 4c. We observe that the three methods are quite competitive. Nevertheless, as shown in Figure 4d, our method displays a better sensitivity in term of classification threshold, since its accuracy outperforms those of the other methods whatever the threshold considered. Moreover, similarly to the ABC class, the F1-score of our method (0.71) is greater than DTFD-MIL (0.67) and CTransPath (0.69) for the GCB class.

4.4 Comparison with medical diagnosis techniques

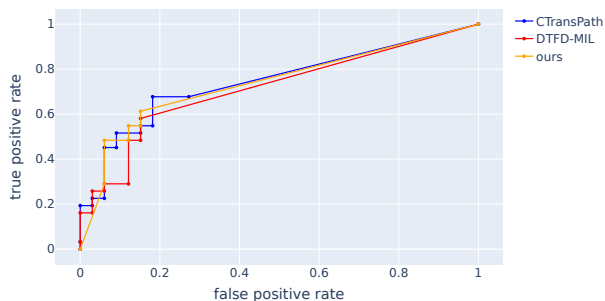
In Table 3, we present a comparative analysis with diagnosis techniques. Here, the student model, trained with our method, is compared with the visual analyses of two pathologists and the Hans algorithm (presented in section 1). On HES images, our model outperforms the two pathologists in distinguishing between the two subtypes, achieving an accuracy that is higher by a margin of at least 9%. The relatively low diagnosis accuracy reached by the pathologists is explained by the fact that they performed their analysis on a single image modality namely HES without any prior knowledge on the IHC modalities. In this sense, we observe that when the pathologists exploit the three IHC modalities, via the Hans algorithm, they reach a better diagnosis accuracy of 83%. Nevertheless, as presented in Table 4, our model demonstrates a lower misclassification rate for ABC patients compared to all other diagnostic techniques." This means that a patient with ABC subtype (subtype with the lowest survival chances) is less likely to be misclassified by our model than the Hans algorithm.



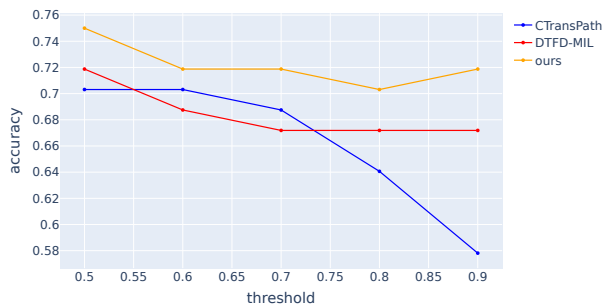
(a) ROC curves of the ABC class.



(b) Accuracy with respect to classifier threshold of the ABC class.



(c) ROC curves of the GCB class.



(d) Accuracy with respect to classifier threshold of the GCB class.

Figure 4: Performance comparisons of state-of-the-art methods obtained on the test set.

We extended our analysis further by visualizing the attention scores of the ViT encoder in the last layer (just before the classification head). Fig. 5 shows results of these scores obtained on two WSI patients. One can observe that areas rich in tumoral cells are receiving high attention.

4.5 Power-law scaling

Scaling law in deep learning refers to the model generalization capability when its computational power and/or training data is scaled [46]. It indicates that increasing resources such as the computing power or volume of training data, can considerably improve the model’s generalization. Studies on ViT [47] have shown that when we increase both the model size and the amount of training data, the generalization error decreases inevitably. In our study, we are interested in observing the student generalization performance when training data is scaled up. In order to achieve this, we trained both teacher and student models on incremental portions of the original dataset, varying from 10% to 100%, measured by number of patients. These subsets

are cumulative, meaning each new subset contains all patients from the previous one, with an added 5% patients of the original training set. Thus, with every step, the subset grows by retaining the existing patients and incorporating new ones. Across all subsets, we train the models with the same initial parameters as well as the same hyperparameters, ensuring equal conditions for every trained models. Observing Fig. 6, the model accuracy exhibits a general increase with the enlargement of the training data. However, distinct declines in accuracy are observed at 90% and 95%, which could be attributed to the incorporation of patient with images markedly different from the existing dataset. Despite this, the introduction of more patients afterwards leads to a rebound in accuracy, highlighting the positive impact on the model’s performance. For instance, the power-law curve shows that the student model could outperform the Hans algorithm and reach an accuracy of 90% with a training set of 327 patients.

Table 3: Comparison of the Student model with clinical diagnosis techniques for DLBCL subtyping on test set (64 patients).

	ABC			GCB			ACC
	PRE	REC	F1	PREC	REC	F1	
pathologist A	0.69	0.67	0.68	0.65	0.65	0.65	0.66
pathologist B	0.57	0.64	0.60	0.54	0.45	0.49	0.55
ours (student)	0.72	0.85	0.78	0.80	0.65	0.71	0.75
IHC (Hans et al. [6])	0.92	0.73	0.81	0.76	0.94	0.84	0.83

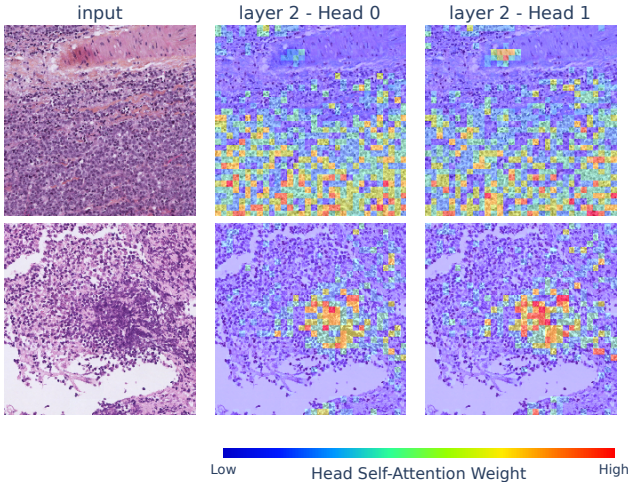


Figure 5: Attention maps generated by our student model on partially tumoral regions.

4.6 Ablation study

This section outlines a series of experiments conducted to evaluate the significance of various features of our architectural design, with results detailed in Table 5. Our exploration initiated with a thorough analysis of KD, a critical component for its role in guiding the student model using insights from the teacher model. We set up an experiment to train the student model exclusively with the true target, without teacher guidance, to evaluate the impact of such self-reliant learning on generalization performance. To assess the effects of independent learning on the student model general-

Table 4: False negative rates.

	ABC	GCB
pathologist A	0.33	0.35
pathologist B	0.36	0.55
ours (student)	0.15	0.35
IHC (Hans et al. [6])	0.27	0.06

ization abilities, we removed the teacher supervision, therefore solely relying the true target one. Our investigation uncovered interesting findings: without teacher guidance the overall accuracy decreased by 5%. This underscores the crucial influence of KD in enhancing the student model’s performance, highlighting the importance of the teacher-student dynamic in the learning process. These insights motivate further research into optimizing KD within our architecture to potentially improve the student model’s performance further.

To minimize the impact of initial parameters choices on the results, we maintained consistent initial model parameters throughout the following experiments. By using the same initial parameters, we could fairly assess the effects of the specific features investigated.

In another experimental setup (*without learnable att.*), substituting learnable attention with dot-product attention led to a notable 17% decrease in accuracy. This resulted in an overall accuracy score even lower than that achieved without any KD at all (58% vs. 70%, respectively). This discrepancy is likely to be attributed to the variations in color distribution between IHC and HES images. These variations may impede the dot-product similarity from effectively capturing relevant information. Therefore, our findings emphasize the crucial role of learnable attention in our methodology. It allows the model to learn to focus on the most relevant IHC information, which is essential for improving the performance of the model.

Next, we removed the softmax activation of the encoder MLP to evaluate its effect on the model’s generalization accuracy (*without softmax*). The modified model demonstrated a substantial decrease in accuracy by 17%, underscoring the essential role of the softmax activation in improving classification performance.

We also investigated how class balancing affects model accuracy (*without class balance*). We trained both the teacher and student models without class balancing. The student model’s accuracy showed a significant decrease of 13%. This suggests that class balancing plays a vital role in the model learning process by ensuring that all classes are equally represented

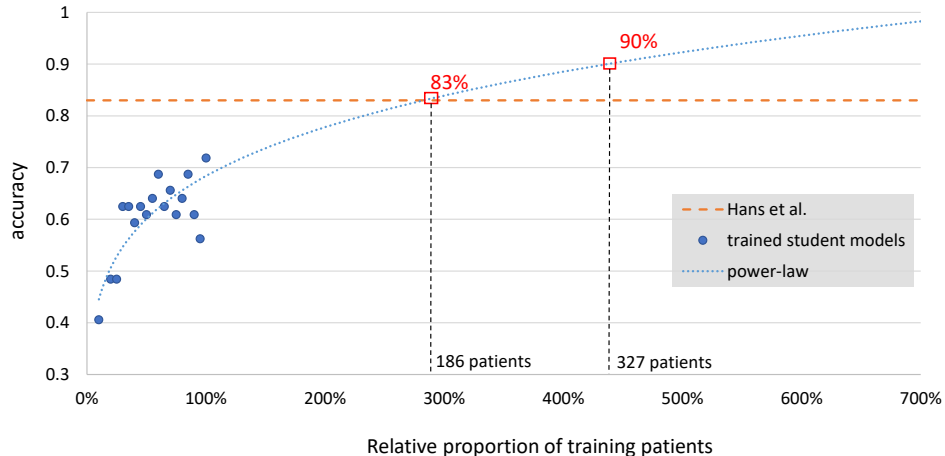


Figure 6: We fit a power-law curve to project the generalization performance of the student model using a larger training set. Accuracy corresponds to the performance of the classification model obtained on the test set of 64 patients.

during training.

Table 5: Ablation results on test set (64 patients). We show the F1-score of each class and the total accuracy.

	F1		ACC
	ABC	GCB	
Original (student) model	0.78	0.71	0.75
without softmax	0.54	0.61	0.58
without class balance	0.62	0.57	0.59
without learnable att.	0.61	0.64	0.63
without KD	0.73	0.67	0.70

4.7 Results on external data

To show the generalization capability of our method, we conduct an experimental study on an external dataset. The dataset is named BCI³ (Breast Cancer Immunohistochemical) [48] and is composed of two modalities (HES and IHC) with 4,873 images each. It is dedicated to the the evaluation of human epidermal growth factor receptor 2 (HER2) expression, an essential factor to formulate a precise treatment for breast cancer [48]. The dataset offers the ground-truth about the expression level of the HER2 factor for each pair of images. In this frame, two recent methods [49, 50] of the state-of-the-art have been proposed to address the classification task of the HES images among four different levels of expression. The dataset meets our

³available at <https://bupt-ai-cz.github.io/BCI/>

needs in term of multiple modalities, which can be exploited by our approach to transfer the knowledge from a HES+IHC classifier to a HES one. However, the images provided in the dataset are not high-resolution WSI as in our case, but rather images of 1024x1024 pixels representing a tiny region of the patient tissue. Therefore, the Vision Transformer-based backbone of our architecture, is not suited for analyzing such form of data, in the sense that the backbone requires in input a sequence of patches uniformly sampled from a large region of tissue of the WSI. Hence, to adapt our architecture to the BCI dataset, we have set up a variant in which we replaced the Transformer-based backbone by a CNN-based one while preserving the core components of our approach, which are the Knowledge Distillation and the multi-modal features fusion mechanisms.

For the choice of the CNN backbone, we exploited the two most competitive methods on our original dataset, namely DTFD-MIL and CTransPath (see Table 2), which permitted to produce three models. To analyze the performance of these models, we also considered in our experimental study the remaining methods of the Table 2; except HIPT and MS-DA-MI, for which their input requirements are not compliant with the BCI dataset. All the methods have been trained accordingly to the settings reported in their original papers. In addition, to highlight the competitiveness of our approach, we also included the results of the state-of-the-art methods (HAHNet [50] and HE-HER2Net [49]) that have been specifically developed for the BCI dataset.

Table 6 summarizes the results obtained on the same test set. The table shows that our three variants im-

prove the accuracy of the original methods with a margin ranging from 3% to 8%. Moreover, we can observe that our variant (based on CTransPath) and the HAHNet model from the state-of-the-art obtained the best accuracy (0.93).

Table 6: Comparison of original state-of-the-art methods and our variant on the BCI [48] test set.

		Accuracy
DTFD-MIL (VGG-19) [13]	original	0.71
	our variant	0.74
DTFD-MIL (ResNet-50) [13]	original	0.83
	our variant	0.86
CTransPath [45]	original	0.85
	our variant	0.93
RSP+CR [43]	–	0.79
KimiaNet [17]	–	0.82
HE-HER2Net [49]	–	0.87
HAHNet [50]	–	0.93

5 Conclusion

We proposed a new methodology based on knowledge distillation to distill a multi-modal vision transformer-based classifier into a mono-modal one for DLBCL molecular subtyping on WSI. We demonstrated the effectiveness of our methodology through an experimental study on 157 patients. Our mono-modal model outperformed state-of-the-art methods, demonstrating the benefit of leveraging multi-modal model knowledge during training. Furthermore, this benefit has been confirmed through an additional experimental study on a external breast cancer dataset (BCI dataset [48]). Inspired by power-law scaling in Vision Transformers, we also explored how our mono-modal model’s diagnostic accuracy scales within larger training data scenarios. Our findings indicate that, with a reasonable number of additional patients, our model can compete with standard clinical methods like IHC [6]. Because of the difficulty of DLBCL subtyping on HES slides, for the same patient, pathologists often rely on several slides to make a diagnosis. Thus, our methodology could be extended to incorporate these HES slides in the student model training to increase the training data. Finally, active learning-driven data selection has been shown to enhance the generalization of deep learning models while minimizing the volume of data required. It could be leveraged to acquire more data efficiently with the objective of reaching competitive performance with the Hans algorithm.

Data Protection

The material for conducting the experiments of this work was declared to the Data Protection Officer (DPO) of CHU de Lille and the National Commission on Informatics and Liberty (CNIL) (declaration number 1129). Furthermore, all the patients were anonymized.

References

- [1] H. Sung, J. Ferlay, R. L. Siegel, M. Laversanne, I. Soerjomataram, A. Jemal, and F. Bray, “Global cancer statistics 2020: Globocan estimates of incidence and mortality worldwide for 36 cancers in 185 countries,” *CA: A Cancer Journal for Clinicians*, vol. 71, pp. 209 – 249, 2021.
- [2] G. Kanas, W. Ge, R. G. W. Quek, K. Keeven, K. Nersesyan, and J. Arnason, “Epidemiology of diffuse large b-cell lymphoma (dlbcl) and follicular lymphoma (fl) in the united states and western europe: population-level projections for 2020–2025,” *Leukemia & Lymphoma*, vol. 63, pp. 54 – 63, 2021.
- [3] S. Mareschal, P. Ruminy, C. Bagacean, V. Marchand, M. Cornic, J. P. Jais, M. Figeac, J. M. Picquenot, T. J. Molina, T. Fest, G. A. Salles, C. Haioun, K. Leroy, H. Tilly, and F. Jardin, “Accurate classification of germinal center b-cell-like/activated b-cell-like diffuse large b-cell lymphoma using a simple and rapid reverse transcriptase-multiplex ligation-dependent probe amplification assay: A calym study.,” *The Journal of molecular diagnostics : JMD*, 2015.
- [4] G. S. Nowakowski, B. R. Laplant, W. R. Macon, C. B. Reeder, J. M. Foran, G. D. Nelson, C. A. Thompson, C. E. Rivera, D. J. Inwards, I. N. Micallef, P. B. Johnston, L. F. Porrata, S. M. Ansell, R. D. Gascoyne, T. M. Habermann, and T. E. Witzig, “Lenalidomide combined with r-chop overcomes negative prognostic impact of non-germinal center b-cell phenotype in newly diagnosed diffuse large b-cell lymphoma: a phase ii study.,” *Journal of clinical oncology : official journal of the American Society of Clinical Oncology*, vol. 33 3, pp. 251–7, 2015.
- [5] G. S. Nowakowski, A. Chiappella, R. D. Gascoyne, D. W. Scott, Q. Zhang, W. Jurczak, M. Özcan, X. nan Hong, J. Zhu, J. Jin, D. Belada, J. M. Bergua, F. Piazza, H. Mociková, A. L. Molinari, D. H. Yoon, F. Cavallo, M. Tani, K. Yamamoto, K. Izutsu, K. Kato, M. S. Czuczman,

- S. Hersey, A. Kilcoyne, J. Russo, K. Hudak, J. Zhang, S. Wade, T. E. Witzig, and U. Vitolo, "Robust: A phase iii study of lenalidomide plus r-chop versus placebo plus r-chop in previously untreated patients with abc-type diffuse large b-cell lymphoma," *Journal of Clinical Oncology*, vol. 39, pp. 1317 – 1328, 2021.
- [6] C. P. Hans, D. D. Weisenburger, T. C. Greiner, R. D. Gascoyne, J. Delabie, G. Ott, H. K. Muller-Hermelink, E. Campo, R. M. Braziel, E. S. Jaffe, Z. Pan, P. Farinha, L. M. Smith, B. Falini, A. H. Banham, A. Rosenwald, L. M. Staudt, J. M. Connors, J. O. Armitage, and W. C. Chan, "Confirmation of the molecular classification of diffuse large b-cell lymphoma by immunohistochemistry using a tissue microarray.," *Blood*, vol. 103 1, pp. 275–82, 2004.
- [7] P. K. Mishro, S. Agrawal, R. Panda, and A. Abraham, "A survey on state-of-the-art denoising techniques for brain magnetic resonance images.," *IEEE reviews in biomedical engineering*, vol. PP, 2021.
- [8] F. Xing, Y. Xie, H. Su, F. Liu, and L. Yang, "Deep learning in microscopy image analysis: A survey," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 29, pp. 4550–4568, 2018.
- [9] M. Z. Atwany, A. Sahyoun, and M. Yaqub, "Deep learning techniques for diabetic retinopathy classification: A survey," *IEEE Access*, vol. PP, pp. 1–1, 2022.
- [10] C. L. Srinidhi, O. Ciga, and A. L. Martel, "Deep neural network models for computational histopathology: A survey," *Medical image analysis*, vol. 67, p. 101813, 2019.
- [11] M. Y. Lu, D. F. K. Williamson, T. Y. Chen, R. J. Chen, M. Barbieri, and F. Mahmood, "Data-efficient and weakly supervised computational pathology on whole-slide images," *Nature Biomedical Engineering*, vol. 5, pp. 555 – 570, 2020.
- [12] X. Wang, H. Chen, C. Gan, H. Lin, Q. Dou, E. Tsougenis, Q. Huang, M. Cai, and P.-A. Heng, "Weakly supervised deep learning for whole slide lung cancer image analysis," *IEEE Transactions on Cybernetics*, vol. 50, pp. 3950–3962, 2020.
- [13] H. Zhang, Y. Meng, Y. Zhao, Y. Qiao, X. Yang, S. E. Coupland, and Y. Zheng, "Dtf-d-mil: Double-tier feature distillation multiple instance learning for histopathology whole slide image classification," *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 18780–18790, 2022.
- [14] S. Graham, Q. D. Vu, S. E. A. Raza, A. Azam, Y.-W. Tsang, J. T. Kwak, and N. M. Rajpoot, "Hover-net: Simultaneous segmentation and classification of nuclei in multi-tissue histology images," *Medical image analysis*, vol. 58, p. 101563, 2018.
- [15] N. Hashimoto, Y. Takagi, H. Masuda, H. Miyoshi, K. Kohno, M. Nagaishi, K. Sato, M. Takeuchi, T. Furuta, K. Kawamoto, K. Yamada, M. Moritsubo, K. Inoue, Y. Shimasaki, Y. Ogura, T. Imamoto, T. Mishina, K. Ohshima, H. Hon-tani, and I. Takeuchi, "Case-based similar image retrieval for weakly annotated large histopathological images of malignant lymphoma using deep metric learning," *Medical image analysis*, vol. 85, p. 102752, 2021.
- [16] N. Hegde, J. D. Hipp, Y. Liu, M. R. Emmert-Buck, E. Reif, D. Smilkov, M. Terry, C. J. Cai, M. B. Amin, C. H. Mermel, P. Q. Nelson, L. H. Peng, G. S. Corrado, and M. C. Stumpe, "Similar image search for histopathology: Smily," *NPJ Digital Medicine*, vol. 2, 2019.
- [17] A. Riasatian, M. Babaie, D. Maleki, S. Kalra, M. Valipour, S. Hemati, M. Zaveri, A. Safarpour, S. Shafiei, M. Afshari, M. Rasooljafari, M. Sikaroudi, M. Adnan, S. Shah, C. Choi, S. Damaskinos, C. J. V. Campbell, P. Diamandis, L. X. Pantanowitz, H. Kashani, A. Ghodsi, and H. R. Tizhoosh, "Fine-tuning and training of densenet for histopathology image representation using tcga diagnostic slides," *Medical image analysis*, vol. 70, p. 102032, 2021.
- [18] H. E. Achi, T. Belousova, L. Chen, A. Wahed, I. Wang, Z. J. Hu, Z. Kanaan, A. Rios, and A. N. D. Nguyen, "Automated diagnosis of lymphoma with digital pathology images using deep learning," *Annals of clinical and laboratory science*, vol. 49 2, pp. 153–160, 2018.
- [19] H. Miyoshi, K. Sato, Y. Kabeya, S. Yonezawa, H. Nakano, Y. Takeuchi, I. Ozawa, S. Higo, E. Yanagida, K. Yamada, K. Kohno, T. Furuta, H. Muta, M. Takeuchi, Y. Sasaki, T. Yoshimura, K. Matsuda, R. Muto, M. Moritsubo, K. Inoue, T. Suzuki, H. Sekinaga, and K. Ohshima, "Deep learning shows the capability of high-level computer-aided diagnosis in malignant lymphoma," *Laboratory Investigation*, pp. 1–11, 2020.

- [20] G. Steinbuss, M. Kriegsmann, C. Zgorzelski, A. Brobeil, B. Goepfert, S. Dietrich, G. Mechtersheimer, and K. Kriegsmann, “Deep learning for the classification of non-hodgkin lymphoma on histopathological images,” *Cancers*, vol. 13, 2021.
- [21] F. Gaudio, R. Tamma, G. Ingravallo, T. Perrone, F. E. Laddaga, M. S. D. Candia, E. Maiorano, D. Ribatti, and G. Specchia, “Computer-driven quantitative image analysis in the assessment of tumor cell and t cell features in diffuse large b cell lymphomas,” *Annals of Hematology*, vol. 97, pp. 663–668, 2018.
- [22] N. Hashimoto, K. Ko, T. Yokota, K. Kohno, M. Nakaguro, S. Nakamura, I. Takeuchi, and H. Hontani, “Subtype classification of malignant lymphoma using immunohistochemical staining pattern,” *International Journal of Computer Assisted Radiology and Surgery*, vol. 17, pp. 1379 – 1389, 2022.
- [23] D. Li, J. R. Bledsoe, Y. Zeng, W. Liu, Y. Hu, K. Bi, A. Liang, and S. Li, “A deep learning diagnostic platform for diffuse large b-cell lymphoma with high accuracy across multiple hospitals,” *Nature Communications*, vol. 11, 2020.
- [24] D. Guidolin, R. Tamma, T. Annese, C. Tortorella, G. Ingravallo, F. Gaudio, T. Perrone, P. Musto, G. Specchia, and D. Ribatti, “Different spatial distribution of inflammatory cells in the tumor microenvironment of abc and gbc subgroups of diffuse large b cell lymphoma,” *Clinical and Experimental Medicine*, vol. 21, pp. 573 – 578, 2021.
- [25] S. Deng, X. Zhang, W. Yan, E. I.-C. Chang, Y. Fan, M. Lai, and Y. Xu, “Deep learning in digital pathology image analysis: a survey,” *Frontiers of Medicine*, vol. 14, pp. 470 – 487, 2020.
- [26] G. Campanella, M. G. Hanna, L. Geneslaw, A. P. Miralflor, V. W. K. Silva, K. J. Busam, E. Brogi, V. E. Reuter, D. S. Klimstra, and T. J. Fuchs, “Clinical-grade computational pathology using weakly supervised deep learning on whole slide images,” *Nature Medicine*, pp. 1–9, 2019.
- [27] R. J. Chen and R. G. Krishnan, “Self-supervised vision transformers learn visual concepts in histopathology,” *ArXiv*, vol. abs/2203.00585, 2022.
- [28] R. J. Chen, C. Chen, Y. Li, T. Y. Chen, A. D. Trister, R. G. Krishnan, and F. Mahmood, “Scaling vision transformers to gigapixel images via hierarchical self-supervised learning,” *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 16123–16134, 2022.
- [29] Z. Shao, H. Bian, Y. Chen, Y. Wang, J. Zhang, X. Ji, and Y. Zhang, “Transmil: Transformer based correlated multiple instance learning for whole slide image classification,” in *Neural Information Processing Systems*, 2021.
- [30] R. Nakhli, P. A. Moghadam, H. Mi, H. S. Farahani, A. S. Baras, C. B. Gilks, and A. Bashashati, “Amigo: Sparse multi-modal graph transformer with shared-context processing for representation learning of giga-pixel images,” *ArXiv*, vol. abs/2303.00865, 2023.
- [31] B. Korbar, A. M. Olofson, A. P. Miralflor, K. M. Nicka, M. A. Suriawinata, L. Torresani, A. A. Suriawinata, and S. Hassanpour, “Deep learning for classification of colorectal polyps on whole-slide images,” *Journal of Pathology Informatics*, vol. 8, 2017.
- [32] M. Kim, H. Sekiya, G. Yao, N. B. Martin, M. Castanedes-Casey, D. W. Dickson, T. H. Hwang, and S. Koga, “Diagnosis of alzheimer’s disease and tauopathies on whole slide histopathology images using a weakly supervised deep learning algorithm,” *Laboratory investigation; a journal of technical methods and pathology*, p. 100127, 2023.
- [33] Y. Dong, Z. Jiang, H. Shen, W. D. Pan, L. A. Williams, V. V. B. Reddy, W. H. Benjamin, and A. W. Bryan, “Evaluations of deep convolutional neural networks for automatic identification of malaria infected cells,” *2017 IEEE EMBS International Conference on Biomedical & Health Informatics (BHI)*, pp. 101–104, 2017.
- [34] C. Dwivedi, S. Nofallah, M. Pouryahya, J. Iyer, K. Leidal, C. Chung, T. Watkins, A. N. Billin, R. Myers, J. Abel, and A. Behrooz, “Multi stain graph fusion for multimodal integration in pathology,” *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 1834–1844, 2022.
- [35] M. Hu, M. Maillard, Y. Zhang, T. Ciceri, G. L. Barbera, I. Bloch, and P. Gori, “Knowledge distillation from multi-modal to mono-modal segmentation networks,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2020.
- [36] Z. Gao, J. Chung, M. Abdelrazek, S. Leung, W. K. Hau, Z. Xian, H. Zhang, and S. Li, “Privileged

- modality distillation for vessel border detection in intracoronary imaging,” *IEEE Transactions on Medical Imaging*, vol. 39, pp. 1524–1534, 2019.
- [37] C. Chen, Q. Dou, Y. Jin, Q. Liu, and P.-A. Heng, “Learning with privileged multimodal knowledge for unimodal segmentation,” *IEEE Transactions on Medical Imaging*, vol. 41, pp. 621–632, 2021.
- [38] S. Wang, Z. Yan, D. Zhang, H. Wei, Z. Li, and R. Li, “Prototype knowledge distillation for medical segmentation with missing modality,” *ArXiv*, vol. abs/2303.09830, 2023.
- [39] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, “An image is worth 16x16 words: Transformers for image recognition at scale,” *ArXiv*, vol. abs/2010.11929, 2020.
- [40] J. Ba, J. R. Kiros, and G. E. Hinton, “Layer normalization,” *ArXiv*, vol. abs/1607.06450, 2016.
- [41] R. Müller, S. Kornblith, and G. E. Hinton, “When does label smoothing help?,” in *Neural Information Processing Systems*, 2019.
- [42] V. Bobée, P. Ruminy, V. Marchand, P.-J. Vially, A. A. Sater, L. Veresezan, F. Drieux, C. Bérard, É. Bohers, S. Mareschal, S. Dubois, J. P. Jais, K. Leroy, M. Figeac, J. M. Picquenot, T. J. Molina, G. A. Salles, C. Haioun, H. Tilly, and F. Jardin, “Determination of molecular subtypes of diffuse large b-cell lymphoma using a reverse transcriptase multiplex ligation-dependent probe amplification classifier: A calym study.,” *The Journal of molecular diagnostics : JMD*, vol. 19 6, pp. 892–904, 2017.
- [43] C. L. Srinidhi, S. W. Kim, F.-D. Chen, and A. L. Martel, “Self-supervised driven consistency training for annotation efficient histopathology image analysis,” *Medical image analysis*, vol. 75, p. 102256, 2021.
- [44] N. Hashimoto, D. Fukushima, R. Koga, Y. Takagi, K. Ko, K. Kohno, M. Nakaguro, S. Nakamura, H. Hontani, and I. Takeuchi, “Multi-scale domain-adversarial multiple-instance cnn for cancer subtype classification with unannotated histopathological images,” *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3851–3860, 2020.
- [45] X. Wang, S. Yang, J. Zhang, M. Wang, J. Zhang, W. Yang, J. Huang, and X. Han, “Transformer-based unsupervised contrastive learning for histopathological image classification,” *Medical image analysis*, vol. 81, p. 102559, 2022.
- [46] S. Abnar, M. Dehghani, B. Neyshabur, and H. Sedghi, “Exploring the limits of large scale pre-training,” *ArXiv*, vol. abs/2110.02095, 2021.
- [47] X. Zhai, A. Kolesnikov, N. Houlsby, and L. Beyer, “Scaling vision transformers,” *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1204–1213, 2021.
- [48] S. Liu, C. Zhu, F. Xu, X. Jia, Z. Shi, and M. Jin, “Bci: Breast cancer immunohistochemical image generation through pyramid pix2pix,” *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 1814–1823, 2022.
- [49] M. S. H. Shovon, M. J. Islam, M. N. A. K. Nabil, M. M. I. Molla, A. I. Jony, and M. F. Mridha, “Strategies for enhancing the multi-stage classification performances of her2 breast cancer from hematoxylin and eosin images,” *Diagnostics*, vol. 12, 2022.
- [50] J. Wang, X. Zhu, K. Chen, L. Hao, and Y. Liu, “Hahnet: a convolutional neural network for her2 status classification of breast cancer,” *BMC Bioinformatics*, vol. 24, 2023.