PCR-MIL: Phenotype Clustering Reinforced Multiple Instance Learning for Whole Slide Image Classification

Jingjiao Lou[1,2], Qingtao Pan[1,2], Qing Yang[3], and Bing Ji[1,2(✉)]

[1] School of Control Science and Engineering, Shandong University, Jinan, China
[2] Key Laboratory of Machine Intelligence and System Control, Ministry of Education, China
b.ji@sdu.edu.cn
[3] Department of Breast and Thyroid Surgery, Shandong Provincial Hospital affiliated to Shandong First Medical University, Jinan, China

Abstract. Multiple instance learning (MIL) has proven effective in classifying whole slide images (WSIs), owing to its weakly supervised learning framework. However, existing MIL methods still face challenges, particularly over-fitting due to small sample sizes or limited WSIs (bags). Pseudo-bags enhance MIL's classification performance by increasing the number of training bags. However, these methods struggle with noisy labels, as positive patches often occupy small portions of tissue, and pseudo-bags are typically generated by random splitting. Additionally, they face difficulties with non-discriminative instance embeddings due to the lack of domain-specific feature extractors. To address these limitations, we propose Phenotype Clustering Reinforced Multiple Instance Learning (PCR-MIL), a novel MIL framework that integrates clustering-based pseudo-bags to improve MIL's noise robustness and the discriminative power of instance embeddings. PCR-MIL introduces two key innovations: (i) Phenotype Clustering-based Feature Selection (PCFS) selects relevant instance embeddings for prediction. It clusters instances into phenotype-specific groups, assigns positive instances to each pseudo-bag, and then uses Grad-CAM to select the most relevant positive embeddings. This approach mitigates noisy label challenges and enhances MIL's robustness to noise; (ii) Reinforced Feature Extractor (RFE) uses reinforcement learning to train an extractor based on selected clean pseudo-bags instead of noisy samples. This approach improves the discriminative power of extracted instance embeddings and enhances the feature representation capabilities of MIL. Experimental results on the publicly available BRACS and CRC-DX datasets demonstrate that PCR-MIL outperforms state-of-the-art methods. The code is available at: https://github.com/JingjiaoLou/PCR-MIL.

Keywords: Whole slide image · Histopathology image classification · Multiple instance learning · Reinforcement learning · Deep learning.

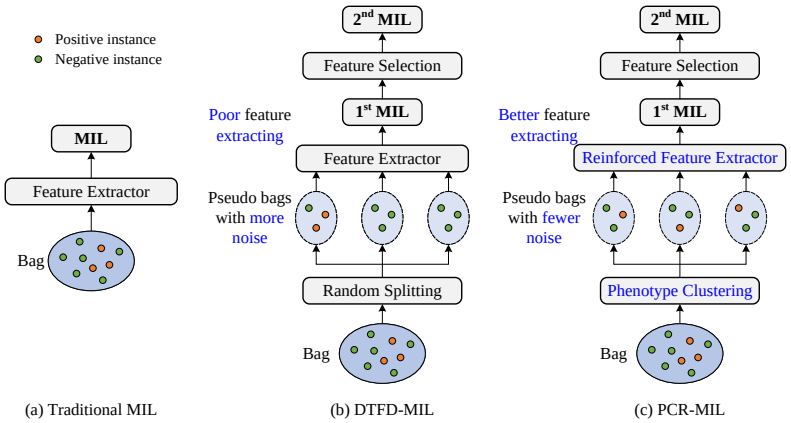

Fig. 1. (a) Traditional MIL takes a WSI bag as input. (b) DTFD-MIL mitigates the overfitting problem caused by a limited number of bags by introducing pseudo-bags. (c) Our proposed PCR-MIL not only addresses the overfitting problem but also enhances MIL's robustness to noise and improves its feature representation capabilities.

1 Introduction

Although multiple instance learning (MIL) has demonstrated strong performance in whole slide image (WSI) classification [4,11,7,5,15,13,8], it still faces the over-fitting problem, where neural network models tend to converge to local minima during optimization [14,16]. The lack of large-scale datasets is one of the most significant factors contributing to this overfitting issue. Despite the enormous size of WSIs, which range from 100 million pixels to 10 gigapixels, each WSI is treated as a single bag under the MIL framework. Pseudo-bags [1,16] offer a potential solution to mitigate the limited number of WSIs by randomly splitting the instances (patches) of a bag (slide) into several smaller bags (pseudo-bags).

However, two challenges prevent MIL with pseudo-bags from handling the over-fitting problem. **i) Pseudo-bags are affected by noisy labels because positive patches occupy only small portions of the tissue, and pseudo-bags are generated through random splitting. (Fig. 2 (a)).** In many histopathology slides, the positive regions corresponding to diseased areas occupy only small portions of the tissue, resulting in a low ratio of positive instances within a slide. As a result, noisy labels emerge when positive instances are not properly split into pseudo-bags. **ii) Embedding-based MIL suffers from non-discriminative features due to the absence of a domain-specific feature extractor. (Fig. 2 (b)).** The embedding-space paradigm has become the primary method in MIL research, heavily relying on the use of an extractor to embed each instance into low-dimensional features, which are then combined to obtain a bag representation. Existing approaches typically use pre-trained models trained on large datasets, such as ImageNet [2], which lack domain-specific knowledge. Although domain-specific extractors are expected to perform

(a) **Challenge 1:** Pseudo-bags suffer from noisy labels because positive patches occupy small portions of tissue, and pseudo-bags are generated through random splitting.

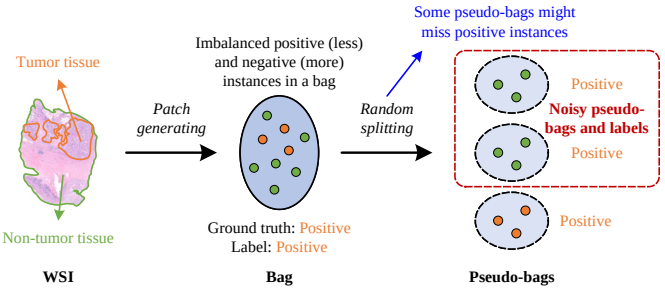

(b) **Challenge 2:** Embedding-based MIL suffers from non-discriminative features due to the lack of a domain-specific feature extractor.

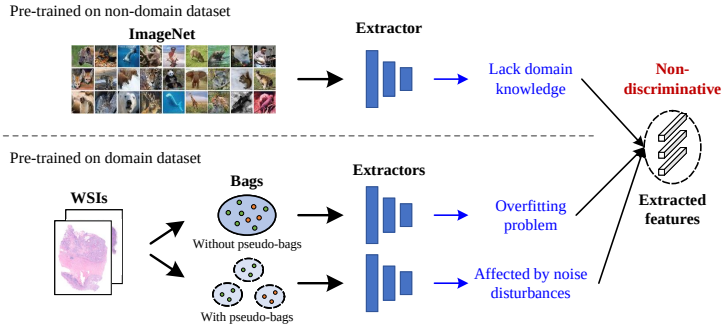

Fig. 2. MIL based on pseudo-bags faces two main challenges for WSI classification: (a) Pseudo-bags suffer from noisy labels because positive patches occupy only small portions of the tissue, and pseudo-bags are generated through random splitting; (b) Embedding-based MIL struggles with non-discriminative features due to the absence of a domain-specific feature extractor.

better, they are hindered by the lack of large-scale datasets and high-quality annotations.

To address the challenges, two novel solutions are proposed (Fig. 3). **i) Our proposed phenotype clustering-based feature selection (PCFS) method addresses noisy labels by splitting positive patches into pseudo-bags and selecting the relevant positive features (Fig. 3 (a)).** PCFS first employs patches from different phenotype clusters to construct pseudo-bags, which contain both positive and negative instances. Then, PCFS selects the instance embeddings with the maximum probability from each pseudo-bag by applying the Grad-based Class Activation Map (Grad-CAM) [12] to identify the signal strengths of instances being positive. This approach mitigates the challenges of noisy labels and enhances the robustness of MIL to noise. **ii) Our proposed Reinforced Feature Extractor (RFE) tackles non-discriminative features by training an optimized extractor on noise-reduced pseudo-bags (Fig. 3 (b)).** RFE selects pseudo-bags with clean (correct) labels to pre-

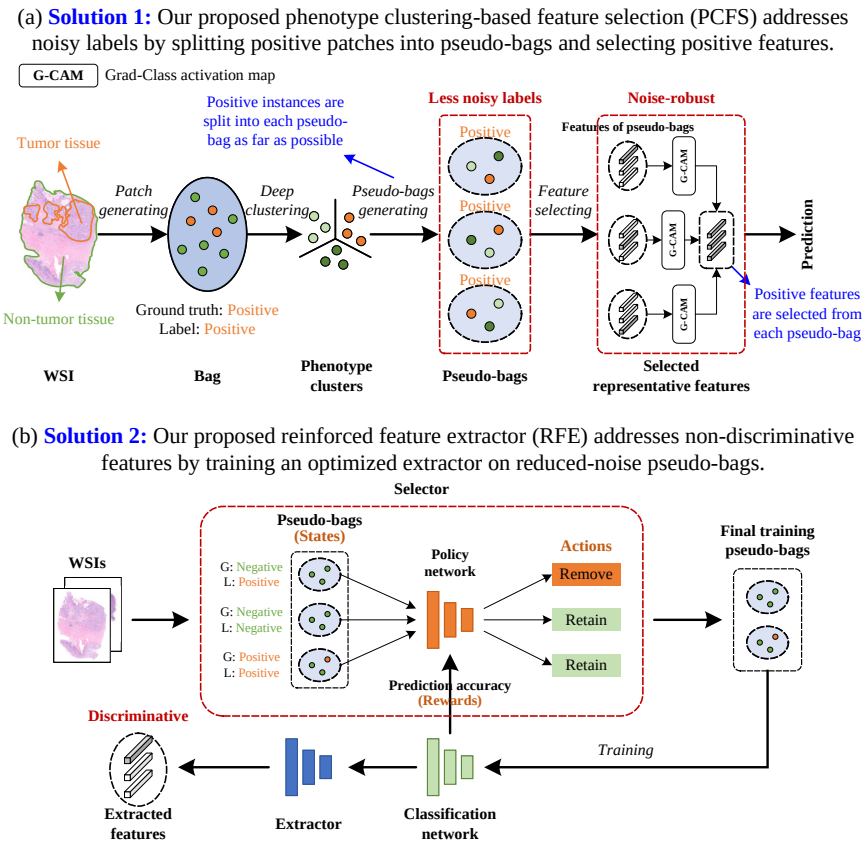

Fig. 3. Our proposed (a) PCFS and (b) RFE address the challenges by selecting the relevant positive embeddings based on phenotype-specific groups and selectively using noisy-reduced pseudo-bags to train a effective domain-specific extractor.

train an extractor by utilizing a policy network from reinforcement learning (RL) [3,9] as a selector to decide whether to retain or remove a pseudo-bag for training. This approach enhances the discriminative power of the extracted instance embeddings and improves MIL's feature representation capabilities.

In this study, we propose a novel double-tier MIL framework specifically designed for WSI classification, which utilizes pseudo-bags for augmentation and a reinforced feature extractor to generate discriminative instance embeddings. The main contributions of this study are:

(1) For the first time, we use phenotype clustering to assign positive instances to each pseudo-bag and further select instance embeddings with strong positive signal strength. This approach mitigates noisy labels and enhances MIL's robustness to noise.

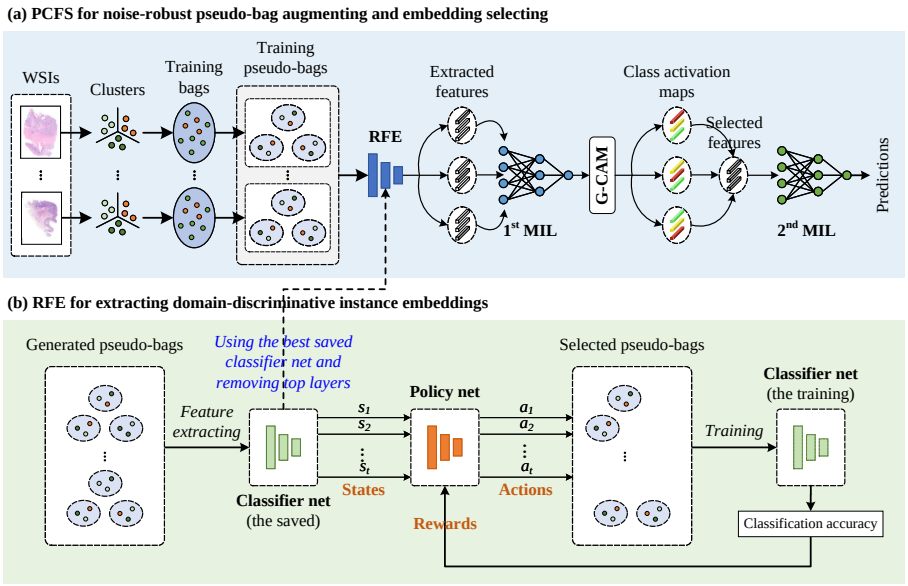

Fig. 4. An overview of our proposed PCR-MIL. PCR-MIL comprises two primary components: (a) PCFS and (b) RFE. PCFS selects representative instance embeddings for final prediction (the 2nd MIL). RFE uses RL to train an extractor to initially extract discriminative instance embeddings (the 1st MIL).

(2) For the first time, we integrate RL into MIL to identify clean pseudo-bags for pre-training a domain-specific feature extractor. This improves the discriminative power of instance embeddings and enhances MIL's feature representation capabilities.

(3) Extensive experimental results on two publicly available datasets demonstrate the superior performance of the proposed method compared to state-of-the-art techniques.

2 Methodology

The PCR-MIL consists of two main components (Fig. 4): (a) PCFS and (b) RFE. Specifically, PCFS selects representative positive instance embeddings for the final prediction, while RFE employs a policy network of RL to train an extractor that initially extracts discriminative instance embeddings.

2.1 PCFS for noise-robust pseudo-bag augmenting and embedding selecting

The proposed PCFS uses patches from different phenotype clusters to construct pseudo-bags and selects instance embeddings with a strong positive signal for

prediction (Fig. 4(a)). Given multiple WSIs from a patient or a single WSI, PCFS divides them into thousands of small patches and applies color normalization using the Macenko method. Specifically, a ResNet50 model pre-trained on ImageNet is used to convert these patches into one-dimensional feature vectors. K-means clustering is then applied to divide these patches into several different clusters based on their corresponding feature vectors. PCFS randomly selects a patch from each cluster to form a pseudo-bag.

PCFS applies the Grad-CAM mechanism to the first MIL to directly infer the signal strength of an instance belonging to a certain class. It then selects the representative instance embeddings based on the signal strength for instance k being class c (where c=0 for negative and c=1 for positive).

Summarized advantages: This approach mitigates the challenges posed by noisy labels and enhances the robustness of MIL to noise.

2.2 RFE for extracting domain-discriminative instance embeddings

The RFE enhances the representation of MIL by training a domain-specific extractor based on noise-reduced pseudo-bags (Fig. 4 (b)). RFE consists of a classifier and a selector. The classifier is built on an attention-based MIL model [5]. It takes a pseudo-bag as input and predicts whether the bag belongs to the positive category. The best pre-trained attention MIL model, excluding the top layers, is later used to extract discriminative instance embeddings.

The selector employs an agent relying on a policy network of reinforcement learning to select pseudo-bags. The state of each time step t is represented as a continuous real-valued vector to meet the Markov decision process. Features extracted by the attention-based MIL are used as states. Each state comprises two types of information: a feature vector of the current bag and the average feature vector of total bags. We define an action $a_i \in (0, 1)$ to determine whether to remove or retain the pseudo-bag. Rewards are derived from the classifier to signify the effectiveness of the selector:

$$R(\tau) = \frac{1}{M} \sum_{t=1}^{M} acc^t - \mu \tag{1}$$

in which τ represents an episode with a sequence of states s and actions a, M represents the total number of pseudo-bags, acc refers to the prediction accuracy used to evaluate the classifier's performance, and μ is a constant.

The policy network optimizes its parameters by maximizing the expected reward. The derivative of the expectation value of $R(\tau)$ is defined as:

$$\nabla \bar{R}_\theta = \frac{1}{N} \sum_{n=1}^{N} R(\tau^n) \nabla log P(\tau^n | \theta) \tag{2}$$

where N represents the total number of episodes, $P(\tau|\theta)$ represents the prediction probability, and θ represents the parameters of the policy network.

Table 1. The comparison results on two datasets show that our PCR-MIL outperforms the state-of-the-art methods. The best values are highlighted in bold.

Methods	BRACS		CRC-DX	
	Accuracy	AUC	Accuracy	AUC
AMIL [5]	73.33	82.89	74.47	73.47
MIL-VT [15]	93.33	98.22	73.40	82.51
TransMIL [13]	80.00	82.89	76.06	75.52
DSMIL [8]	100.00	100.00	79.79	77.56
DTFD-MIL [16]	86.67	93.78	73.40	79.22
PCR-MIL (ours)	**100.00**	**100.00**	**84.57**	**84.12**

Subsequently, the gradient of $\bar{R}_\theta$ can be further expanded to:

$$\nabla \bar{R}_\theta = \frac{1}{N} \sum_{n=1}^{N} \sum_{t=1}^{T} R(\tau^n) \nabla log p(a^t | s^t, \theta) \tag{3}$$

Summarized advantages: This approach enhances the discriminative power of the extracted instance embeddings and improves the feature representation capabilities of MIL.

3 Experiment

3.1 Datasets Description

BRACS The BRACS [10] is a breast cancer dataset, which comprises 4391 regions of interest (ROIs) extracted from 325 hematoxylin and eosin (H&E) breast carcinoma WSIs. In this study, we employ the BRACS dataset for binary classification, differentiating between invasive breast cancer and a combined category of non-invasive cases.

CRC-DX The CRC-DX [6] is a colorectal dataset. It divides 360 patients into training and testing datasets, comprising 39 MSI and 221 MSS, and 26 MSI and 74 MSS, respectively. It comprises 193,312 image patches derived from histological images of CRC patients in the Cancer Genome Atlas (TCGA) cohort.

3.2 Implementation Details

Our code was implemented using TensorFlow and run on the Ubuntu 20.04.4 LTS operating system with Nvidia GeForce RTX 3090 GPUs (24GB). To ensure a fair comparison, the proposed method and other methods are constructed using identical architecture, mini-batch size, and parameter initialization. Each method performed with consistent data splitting.

Table 2. Ablation studies demonstrate the effectiveness of each key component of our proposed PCR-MIL method.

Number	Baseline	PCFS	RFE	BRACS		CRC-DX	
				Accuracy	AUC	Accuracy	AUC
No. 1	✓			86.67	93.78	73.40	79.22
No. 2	✓	✓		93.33	95.11	75.53	78.69
No. 3	✓		✓	100.00	100.00	74.47	80.19
No. 4	✓	✓	✓	**100.00**	**100.00**	**84.57**	**84.12**

3.3 Results and Discussion

Comparison Study The comparison study shows that our PCR-MIL achieves highly competitive performance. Table 1 presents the outcomes of various methods, including attention-based MIL [5], MIL-VT [15], TransMIL [13], DTFDMIL [8] and our proposed PCR-MIL. The performance metrics include Accuracy (%) and AUC (%), as summarized in Table 1. On the BRACS dataset, PCR-MIL achieves perfect scores in both Accuracy (100.00%) and AUC (100.00%), matching the performance of DSMIL while surpassing other state-of-ther-art methods. For instance, MIL-VT attains an AUC of 98.22%, whereas AMIL and TransMIL yield lower results. For the CRC-DX dataset, PCR-MIL demonstrates superior generalization, achieving an Accuracy of 84.57% and AUC of 84.12%, outperforming all competing methods. The second-best performer, DSMIL, attains an Accuracy of 79.79% and AUC of 77.56%, followed by TransMIL (76.06% Accuracy and 75.52% AUC). Notably, PCR-MIL improves Accuracy by 4.78 percentage points and AUC by 6.56 percentage points over DSMIL on CRC-DX.

Ablation Study As demonstrated in Table 2, we systematically evaluate the contributions of key components in our proposed PCR-MIL framework by conducting ablation studies on both BRACS and CRC-DX datasets. The baseline model (DTFD-MIL) achieves an Accuracy of 86.67% and AUC of 93.78% on BRACS, while yielding 73.40% Accuracy and 79.22% AUC on CRC-DX. The results demonstrate that incorporating the PCFS and RFE modules improve models' performance, respectively. The integration of both components achieves state-of-the-art results.

4 Conclusion

In this study, we propose a novel double-tier MIL network, named PCR-MIL, for whole slide image classification. The PCFS innovatively mitigates noise interference by constructing phenotype pseudo-bags and selecting the most representative positive instance embeddings, thereby improving the noise robustness of MIL. The novel RFE trains an effective domain-specific feature extractor by autonomously selecting clean pseudo-bags for training, thereby enhancing the feature representation of MIL. Extensive experimental results demonstrate the effectiveness of both our proposed PCR-MIL method and its individual components.

Acknowledgments. This work was partly supported by the National Natural Science Foundation of China (Grant No.62173212), Taishan Scholars Program of Shandong Province (Grant No.tsqn202306017) and Shandong Province "Double-Hundred Talent Plan" on 100 Foreign Experts and 100 Foreign Expert Teams (Grant No.WSR2023049).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Chen, Y., Shao, Z., Bian, H., Fang, Z., Wang, Y., Cai, Y., Wang, H., Liu, G., Li, X., Zhang, Y.: dmil-transformer: Multiple instance learning via integrating morphological and spatial information for lymph node metastasis classification. IEEE Journal of Biomedical and Health Informatics **27**(9), 4433–4443 (2023)
2. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
3. Fei, H., Zhang, Y., Ren, Y., Ji, D.: Optimizing attention for sequence modeling via reinforcement learning. IEEE Transactions on Neural Networks and Learning Systems **33**(8), 3612–3621 (2021)
4. Gadermayr, M., Tschuchnig, M.: Multiple instance learning for digital pathology: A review of the state-of-the-art, limitations & future potential. Computerized Medical Imaging and Graphics p. 102337 (2024)
5. Ilse, M., Tomczak, J., Welling, M.: Attention-based deep multiple instance learning. In: International conference on machine learning. pp. 2127–2136. PMLR (2018)
6. Kather, J.N., Pearson, A.T., Halama, N., Jäger, D., Krause, J., Loosen, S.H., Marx, A., Boor, P., Tacke, F., Neumann, U.P., et al.: Deep learning can predict microsatellite instability directly from histology in gastrointestinal cancer. Nature medicine **25**(7), 1054–1056 (2019)
7. Van der Laak, J., Litjens, G., Ciompi, F.: Deep learning in histopathology: the path to the clinic. Nature medicine **27**(5), 775–784 (2021)
8. Li, B., Li, Y., Eliceiri, K.W.: Dual-stream multiple instance learning network for whole slide image classification with self-supervised contrastive learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 14318–14328 (2021)
9. Li, H., Xu, H.: Deep reinforcement learning for robust emotional classification in facial expression recognition. Knowledge-Based Systems **204**, 106172 (2020)
10. Pati, P., Jaume, G., Foncubierta-Rodriguez, A., Feroce, F., Anniciello, A.M., Scognamiglio, G., Brancati, N., Fiche, M., Dubruc, E., Riccio, D., et al.: Hierarchical graph representations in digital pathology. Medical image analysis **75**, 102264 (2022)
11. Schirris, Y., Gavves, E., Nederlof, I., Horlings, H.M., Teuwen, J.: Deepsmile: contrastive self-supervised pre-training benefits msi and hrd classification directly from h&e whole-slide images in colorectal and breast cancer. Medical Image Analysis **79**, 102464 (2022)
12. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. In:

Proceedings of the IEEE international conference on computer vision. pp. 618–626 (2017)
13. Shao, Z., Bian, H., Chen, Y., Wang, Y., Zhang, J., Ji, X., et al.: Transmil: Transformer based correlated multiple instance learning for whole slide image classification. Advances in neural information processing systems **34**, 2136–2147 (2021)
14. Song, H., Tang, J., Xiao, H., Hu, J.: Rethinking overfitting of multiple instance learning for whole slide image classification. In: 2023 IEEE International Conference on Multimedia and Expo (ICME). pp. 546–551. IEEE (2023)
15. Yu, S., Ma, K., Bi, Q., Bian, C., Ning, M., He, N., Li, Y., Liu, H., Zheng, Y.: Mil-vt: Multiple instance learning enhanced vision transformer for fundus image classification. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VIII 24. pp. 45–54. Springer (2021)
16. Zhang, H., Meng, Y., Zhao, Y., Qiao, Y., Yang, X., Coupland, S.E., Zheng, Y.: Dtfd-mil: Double-tier feature distillation multiple instance learning for histopathology whole slide image classification. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 18802–18812 (2022)