Reward-Guided Distillation: A Progressive Pseudo-bag Purification Framework for WSI Multiple Instance Learning

Qi Jia[1], Pengshuo Zhang[1], Shuilian Yao[1], Yu Liu[1]*, Lili Sun[2], Yanmei Zhu[2], and Xin Fan[1]

[1] Dalian University of Technology, China
[2] Liaoning Cancer Hospital and Institute, Cancer Hospital of China Medical University, China
liuyu8824@dlut.edu.cn

Abstract. Whole-slide image (WSI) analysis is foundational to digital pathology, yet granular four-class metastasis grading (Normal, ITC, Micro, and Macro-metastasis) remains inherently challenging. Existing pseudo-bag strategies typically aggregate features from all generated bags for classification. However, such stochastic grouping inevitably introduces severe semantic dilution and label noise, especially in sparse categories like ITC and Micro-metastasis, where numerous pseudo-bags may be entirely void of discriminative signals. To mitigate this, we propose a novel tiered MIL framework designed to strategically filter out class-irrelevant noise. Specifically, a Progressive Shapley-based module is introduced to quantify the diagnostic contribution of each patch, facilitating the construction of ordered, high-quality pseudo-bags. Subsequently, a Reward-Guided Selection Module (RGSM) is employed to estimate pseudo-bag labels and dynamically discard those exhibiting conflicts with the global slide-level ground truth. Extensive experiments on a newly curated four-class WSI dataset demonstrate that our purification strategy achieves state-of-the-art performance. Notably, it significantly enhances the detection of highly sparse metastasis, outperforming the best baseline by 10.66% in F1-score. The code will be available at https://github.com/PS-zhang/RGSM

Keywords: Whole Slide Imaging · Multiple Instance Learning · Pseudo-bag Purification

1 Introduction

Whole Slide Image (WSI) analysis is foundational to digital pathology. Due to the gigapixel scale and immense computational overhead of WSIs [11], the standard approach is a "divide-and-conquer" Multiple Instance Learning (MIL) paradigm, which partitions a WSI into thousands of small, localized patches. As WSI classification is inherently a weakly supervised problem—where a single

* Corresponding author.

slide-level label corresponds to a vast bag of instances—the core challenge lies in effectively identifying and aggregating discriminative patches that truly represent the underlying pathology [12].

To bridge the supervision gap, pseudo-bagging has emerged as a critical strategy, where patches are grouped into smaller subsets to refine feature granularity [14, 3, 27, 21, 15, 10]. Early methods relied on stochastic grouping,this random approach forces irrelevant noise into pseudo-bags, severely compromising the representation of minority classes. While recent advances, such as the Shapley-based approach [24], aim to ensure that positive patches are distributed across bags. However, a persistent bottleneck remains: current frameworks typically aggregate all patches within a pseudo-bag during the second stage of feature distillation. This indiscriminately incorporates class-irrelevant instances, introducing significant noise and leading to suboptimal representation [22].

This dilemma is particularly acute in the four-class metastasis grading of lymph nodes. Unlike binary classification which merely detects the presence of malignancy, this task requires a granular assessment of metastatic extent, encompassing Normal, Isolated Tumor Cells (ITC), Micro-metastasis (Micro), and Macro-metastasis (Macro). In these refined categories, particularly ITC, the rarity of malignant patches means that stochastic grouping often produces pseudo-bags where positive signals are overwhelmed by negative instances, or bags that are entirely void of discriminative content. Consequently, traditional feature distillation processes inevitably ingest these irrelevant features, causing severe "semantic dilution" and diagnostic errors [2].

To address these challenges, we propose a tiered MIL framework centered on Pseudo-bag Purification. We argue that the key to robust grading lies in abandoning blind aggregation in favor of a value-driven, denoising representation [7, 9, 16, 19, 28]. Specifically, our framework first employs a Progressive Shapley-based module to construct systematically ordered pseudo-bags based on their diagnostic contribution. Subsequently, we propose a Reward-Guided Selection Module (RGSM) consisting of a pre-trained Generator and a Rewarder. The Generator estimates pseudo-labels for each bag, while the Rewarder calculates a reward score based on the cosine similarity between pseudo-labels and ground-truth slide labels [13]. This smooth, well-calibrated metric allows the framework to dynamically discard class-irrelevant bags, effectively eliminating redundant noise before final classification.

The primary contributions of this work are summarized as follows:

- We propose a pseudo-bag purification framework that strategically filters out irrelevant bag-level features, directly addressing the feature drift caused by noisy instances in multi-class WSI analysis.
- We propose the Reward-Guided Selection Module (RGSM), which leverages a Generator-Rewarder architecture to quantify bag quality. By utilizing a cosine-similarity-based reward score, it isolates high-fidelity signals and prunes non-contributory pseudo-bags.
- Our framework achieves state-of-the-art performance on the newly curated NIMM dataset for four-class metastasis grading. Experimental results demon-

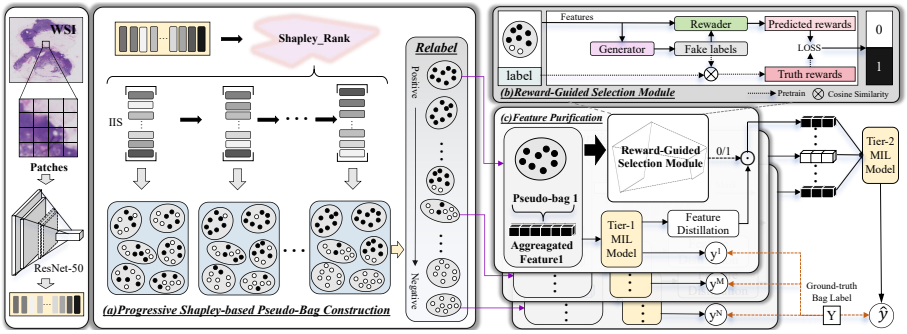

Fig. 1. The overview of our proposed tiered MIL framework. (a) **Progressive Shapley-based Pseudo-Bag Construction**: Guided by Shapley ranks, instances undergo an iterative grouping process and are relabeled into a graded series of pseudo-bags. (b) **Reward-Guided Selection Module**: a generator-rewarder architecture produces a binary mask $M \in \{0, 1\}$ to filter pseudo-bag. (c) **Feature Purification**: the purified features, processed through RGSM, are aggregated by a Tier-2 MIL model for final bag-level prediction $\hat{y}$.

strate superior accuracy and interpretability, particularly in identifying the subtle morphological signatures of ITC and Micro.

2 Methodology

As illustrated in Fig. 1, the proposed framework operates through a tiered purification pipeline. Initially, each whole slide image is partitioned into a myriad of localized patches, from which preliminary features are extracted using a pre-trained ResNet-50 backbone. These features are then evaluated via a Progressive Shapley-based module to quantify the diagnostic contribution of each patch, facilitating the construction of ordered pseudo-bags. As depicted in the progression from left to right in the figure, the ranking mechanism undergoes iterative refinement during training. In the initial stages, malignant instances (denoted by black dots) and negative instances (white dots) may be stochastically distributed. As the model converges, the Shapley-based ranking becomes increasingly precise, effectively clustering malignant patches into specific bags while isolating negative ones. Subsequently, a relabeling strategy is applied to these ordered bags: those with high diagnostic relevance inherit the global slide-level ground truth, while lower-ranked bags are assigned negative labels. To further eliminate class-irrelevant noise, these preliminarily ordered pseudo-bags undergo rigorous screening via the Reward-Guided Selection Module (RGSM), yielding distilled, high-fidelity features. Finally, these purified representations are aggregated by a Tier-2 MIL head to produce the final diagnostic prediction $\hat{y}$.

2.1 Progressive Shapley-based Pseudo-Bag Construction

Initially, the input WSI is partitioned into patches and processed through a ResNet-50 backbone to extract instance-level features $\mathbf{f}_i$. To evaluate instance importance, we adopt a Shapley value-based method [24, 4]. Specifically, this approach begins by randomly grouping instances into initial subsets. By iteratively computing the marginal contribution of each patch to randomly sampled background subsets, it progressively ranks the instances, ultimately generating a well-ordered series of pseudo-bags based on their confidence. Formally, the Shapley score Φ_i for each instance i is estimated as follows:

$$\Phi_i = \text{Norm}\left(\frac{1}{M \cdot T}\sum_{j=1}^{M}\sum_{i=1}^{T}\left[f(S_j \cup \{\mathbf{f}_i\}) - f(S_j)\right]\right), \tag{1}$$

where S_j is a sampled background subset, M and T denote the number of subsets and sampling iterations, and $f(\cdot)$ is the prediction function, mapping subset features to a WSI-level predictive probability. During training, the Shapley values Φ_i iteratively rank instances by diagnostic relevance, ensuring high-scoring patches are preferentially assigned to leading pseudo-bags [23].

Pseudo-bag Label Relabeling To address potential inconsistencies between the generated pseudo-labels and the global ground truth, a relabeling strategy is applied [5]. Specifically, instances in pseudo-bags that exhibit low contribution or label conflict are reassigned as normal samples. Thus we get pseudo-bag labels l. This progressive refinement ensures that the tiered MIL architecture receives more accurate supervision signals as training proceeds.

2.2 Reward-Guided Selection Module

To mitigate inherent label noise in pseudo-bags, we propose the RGSM, which collaboratively optimizes bag-level supervision via a generator and a rewarder. During pre-training, the generator produces fake labels, which the rewarder processes alongside input features to predict rewards. The network is optimized by minimizing the loss between these predictions and ground-truth rewards, the latter calculated via the cosine similarity between the pseudo-bag's given labels and fake labels. During the training phase, the pseudo-bag features alongside their corresponding labels are directly fed into the rewarder for continuous learning. Finally, the input pseudo-bag features are processed by the trained rewarder to produce confidence scores.

Pseudo-label Generator As depicted in the upper panel of Fig. 2, the input pseudo-bag feature $\mathbf{b} \in \mathbb{R}^d$ first passes through a dimensionality reduction module,where d is 2048, which is composed of a sequence of Linear layers interleaved with ReLU activations. Subsequently, a final MLP block processes these reduced representations to generate a continuous output g_{cont} and a discretized fake

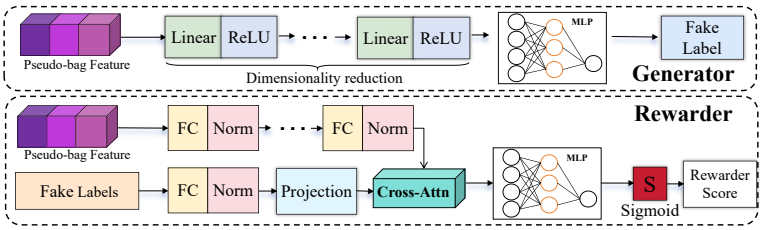

Fig. 2. Detailed architectures of the Generator and Rewarder. The Generator outputs fake labels via dimensionality reduction and MLP, while the Rewarder computes reward scores using cross-attention on pseudo-bag features and fake labels.

label $\hat{g}$. Formally, denoting this entire sequential mapping as $\mathrm{MLP}_G(\mathbf{b};\phi)$, the generation process is formulated as:

$$g_{cont} = \mathrm{ReLU}(\mathrm{MLP}_G(\mathbf{b};\phi)), \quad \hat{g} = \mathrm{clamp}\left(\lfloor g_{cont} \cdot K \rfloor, 0, K-1\right), \tag{2}$$

where ϕ denotes the generator parameters and K is the total number of classes. The floor operation $\lfloor \cdot \rfloor$ is applied to round the scaled value $g_{cont} \cdot K$ down to the nearest integer. Subsequently, the clamp function is employed to strictly bound this integer within the valid range of categorical indices $[0, K-1]$.

Rewarder As illustrated in the lower panel of Fig. 2, the rewarder R, parameterized by θ, is designed to evaluate the diagnostic importance of each pseudo-bag. The rewarder first processes the inputs through two parallel pathways. Specifically, the pseudo-bag feature $\mathbf{b}$ is mapped through a sequence of Fully Connected (FC) layers interleaved with Normalization operations. Concurrently, the fake label $\hat{g}$ is mapped into a rich label embedding $\mathbf{e}_{\hat{g}}$ via a separate branch consisting of FC, Normalization and Projection layers. These refined representations are then fused via a cross-attention mechanism R_{attn}. The cross-attended features are subsequently passed through a MLP block and a Sigmoid activation to yield the predicted reward $\hat{r}$:

$$\hat{r} = \mathrm{Sigmoid}(\mathrm{MLP}(R_{attn}(\mathbf{b}, \mathbf{e}_{\hat{g}}; \theta))). \tag{3}$$

Optimization and Filter The truth rewards R_{gt} are derived by computing the cosine similarity between the fake labels $\hat{g}$ and the pseudo-bag labels l derived from the Shapley stage. Subsequently, the rewarder is optimized by minimizing the mean squared error between the predicted reward $\hat{r}$ and the truth rewards R_{gt} :

$$R_{gt} = \frac{\hat{g} \cdot l}{\|\hat{g}\|_2 \|l\|_2}, \mathcal{L}_R = \left\|\hat{r} - R_{gt}\right\|^2. \tag{4}$$

Based on these scores, a binary mask is applied to filter out noisy pseudo-bags, ultimately yielding purified features for downstream tasks.

2.3 Feature Purification

As shown in Fig. 1, the tier-1 MIL model first processes the n pseudo-bags to output initial predictions $\{y^1, \ldots, y^n\}$. Filtered by the RGSM mask, the retained high-quality pseudo-bags are then fed into the tier-2 MIL model to yield the final prediction $\hat{y}$. The overall classification loss $\mathcal{L}_{cls}$ jointly optimizes both tiers against the ground truth label Y. To train the framework end-to-end, the classification loss $\mathcal{L}_{cls}$ and rewarder loss $\mathcal{L}_R$ are jointly minimized:

$$\mathcal{L}_{cls} = \alpha \frac{1}{N} \sum_{i=1}^{N} \mathcal{L}(y^i, Y) + (1-\alpha)\mathcal{L}(\hat{y}, Y), \quad \mathcal{L}_{total} = \mathcal{L}_{cls} + \lambda \mathcal{L}_R, \tag{5}$$

where α and λ are hyperparameters balancing the loss contributions between the two tiers and the dual objectives, respectively, to progressively enhance pseudo-bag construction and selection.

3 Experiment

3.1 Experiment Settings

NIMM datasets We conduct experiments on the newly constructed NIMM dataset. The name is derived from the initials of the four-class classification: Negative, Isolated Tumor Cells, Micro-metastasis, and Macro-metastasis [1]. All tissue specimens are processed following standard clinical protocols. The final four-class annotation distribution consists of 231 Negative cases, 40 ITCs, 76 Micro-metastasis cases, and 282 Macro-metastasis cases.

Evaluation Metrics The **AUC** assesses the model's capability to discriminate between different classes, with values ranging from 0 to 1. The **F1-score** is the harmonic mean of Precision and Recall, defined as $2 \times \frac{Precision \times Recall}{Precision + Recall}$, where Precision is computed as $\frac{TP}{TP+FP}$ and Recall as $\frac{TP}{TP+FN}$.

Implementation Details Experiments were conducted on two RTX 3060 GPUs for 200 epochs using Adam optimizer with a learing rate of $l0^{-4}$, a weight decay of 10^{-4} and an early-stopping patience of 20 epochs. For Shapley-based construction, we use $M = 10$ subsets, an update step of 2, and a 0.3 relabeling ratio for ITC and Micro. RGSM is pre-trained for 50 epochs. The total objective coefficient α and λ is set to 0.5 and 0.4, respectively.

3.2 Experiment Results

Results on the NIMM Dataset We evaluate our framework against seven MIL methods, strategically selected to cover classic attention mechanisms [8], advanced sequence modeling [18, 25] and recent noise-robust pseudo-bag paradigms [26, 20, 27, 17]. As shown in Table 1, our method achieves state-of-the-art performance

across most metrics. It achieves an absolute improvement of 3.64% in the overall F1-score compared to the second-best result. Specifically, it yields a 1.86% increase in the F1-score of ITC over MRePath.

Table 1. Results are reported as $\text{Mean}_{\text{Standard Deviation}}$. **Bold** and underlined values indicate the best and second-best performance, respectively.

Method	Pub'Year	Overall		ITC		Micro	
		AUC	F1-score	AUC	F1-score	AUC	F1-score
DFTD [26]	CVPR'22	$91.51_{0.49}$	$68.19_{1.18}$	$92.73_{1.62}$	$27.61_{2.95}$	$91.90_{0.50}$	$49.91_{10.58}$
AB-MIL [8]	ICML'18	$96.05_{1.57}$	$69.01_{1.95}$	$\mathbf{95.16}_{2.12}$	$44.44_{8.58}$	$91.62_{3.09}$	$41.75_{2.96}$
TransMIL [18]	Neurips'21	$96.15_{0.69}$	$71.50_{2.46}$	$94.86_{0.66}$	$47.87_{8.00}$	$91.63_{1.57}$	$47.82_{3.93}$
MambaMIL [25]	MICCAI'24	$93.14_{1.89}$	$70.49_{6.30}$	$89.94_{3.41}$	$41.63_{12.88}$	$85.61_{4.20}$	$49.21_{11.97}$
TORCH [20]	Nat. Med.'24	$96.05_{0.29}$	$71.03_{2.36}$	$93.95_{0.85}$	$46.26_{8.41}$	$\underline{92.83}_{0.70}$	$46.46_{3.57}$
PAMOE [27]	CVPR'25	$96.43_{0.67}$	$72.55_{1.09}$	$94.92_{1.36}$	$49.37_{8.19}$	$92.82_{1.99}$	$48.92_{6.06}$
MRePath [17]	IJCAI'25	$\underline{96.21}_{0.58}$	$\underline{74.09}_{0.96}$	$\underline{95.11}_{0.54}$	$\underline{52.69}_{9.60}$	$92.20_{2.15}$	$\underline{51.63}_{4.77}$
Ours	-	$\mathbf{96.58}_{0.37}$	$\mathbf{77.73}_{0.88}$	$91.20_{1.07}$	$\mathbf{54.55}_{5.06}$	$\mathbf{94.82}_{2.21}$	$\mathbf{62.29}_{5.61}$

3.3 Ablation Studies

As presented in Table 2, we select DFTD [26] as our baseline and conduct ablation studies on the RGSM and Shapley modules. Specifically, compared to the baseline, our proposed method achieves an absolute AUC improvement of 4.33%. For the minority class ITC, the RGSM and Shapley modules achieve remarkable F1-score improvements of 4.24% and 19.28%. This demonstrates their strong capability in handling minority classes.

To investigate the impact of the initial number of pseudo-bags on model performance, we present the AUC and F1-score across varying pseudo-bag quantities in Fig. 3(A). The results show that the performance peaks when the number of pseudo-bags is 10. Therefore, we set the pseudo-bag count to 10 for all our experiments.

Table 2. Ablation study of the proposed components on NIMM dataset.

	Method			Overall		ITC		Micro	
	Baseline	RGSM	Shapley	AUC	F1-Score	AUC	F1-Score	AUC	F1-Score
(a)	✓			$91.51_{0.49}$	$68.19_{1.18}$	$92.73_{1.62}$	$27.61_{2.95}$	$93.90_{0.50}$	$51.91_{10.58}$
(b)	✓	✓		$94.99_{0.97}$	$67.96_{0.74}$	$93.17_{0.29}$	$31.85_{2.57}$	$93.48_{0.47}$	$54.36_{3.42}$
(c)	✓		✓	$95.54_{0.33}$	$71.71_{0.66}$	$86.13_{0.82}$	$46.89_{0.88}$	$93.59_{0.26}$	$\mathbf{64.95}_{6.25}$
(d)	✓	✓	✓	$\mathbf{95.84}_{0.37}$	$\mathbf{77.73}_{0.88}$	$\mathbf{91.20}_{1.07}$	$\mathbf{54.55}_{5.06}$	$\mathbf{94.82}_{2.21}$	$62.29_{5.61}$

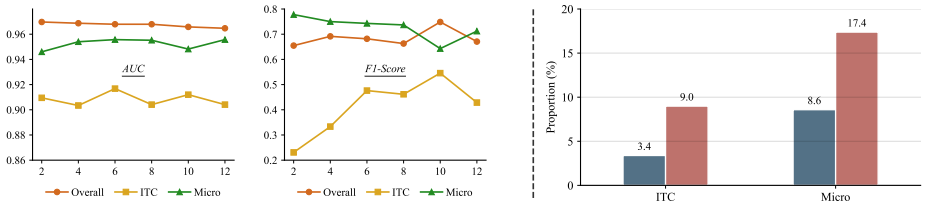

Fig. 3. (A) Performance variations in terms of AUC and F1-score, where the x-axis represents the number of pseudo-bags. (B) Comparative analysis of the cancer patch proportions before and after applying the RGSM filtering module.

3.4 Result Visualization

To quantitatively analyze the specific patches filtered out by the RGSM, we employ **HoVer-Net** [6] for validation. By evaluating the cellular status within each patch, HoVer-Net can reliably derive the patch-level ground truth labels. As illustrated in Fig. 3(B), where the proportion is calculated as the average ratio of class-specific patches across all pseudo-bags, applying the RGSM increases this metric by 5.6% for ITC and 8.8% for Micro. This significant enrichment demonstrates that the RGSM effectively accentuates rare yet diagnostically critical malignant features.

Figure 4 illustrates the trajectory of the patch filtering process guided by RGSM, presented via *t*-SNE visualization of a representative WSI. As observed, the model effectively identifies and eliminates the most typical non-neoplastic patches. Although minor misclassifications inevitably occur, the dense and purified tumor core regions are robustly preserved. This preservation demonstrates the successful extraction of high-quality features.

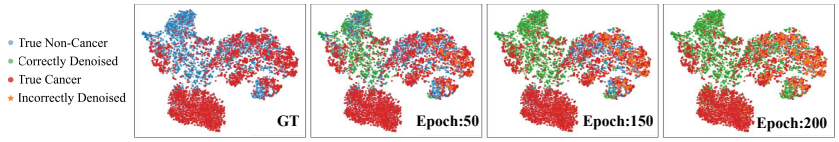

Fig. 4. Visualization of the progressive patch denoising guided by the RGSM. Blue and red dots denote ground-truth non-neoplastic and neoplastic patches, respectively. Green and orange dots represent correctly and incorrectly denoised instances.

4 Conclusion

This paper proposes a feature purification framework to address pseudo-bag quality bottlenecks in Multiple Instance Learning. By integrating a progressive

Shapley-based construction mechanism with RGSM, our approach is the first to employ strategic pseudo-bag filtering. Experimental results on four-class classification tasks demonstrate a 10.66% the F1-score of Micro improvement over state-of-the-art methods. Future research will focus on evaluating the generalizability of this framework across diverse weakly-supervised clinical applications.

Acknowledgments. This work is supported by the National Natural Science Foundation of China (62272083, 62495085, 62472066), the Key Research and Development Program of Liaoning Province (2023JH26/10200014), the Fundamental Research Funds for the Central Universities of China (LD2023019), the Liaoning Provincial Natural Science Foundation of China (2024-MS-271, 2025-MS-336), and the Liaoning Provincial Medical Education Research Project of China (2024-N002-04).

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

References

1. Bándi, P., et al.: From detection of individual metastases to classification of lymph node status at the patient level: the camelyon17 challenge. IEEE Transactions on Medical Imaging **38**(2), 550–560 (2018)
2. Bourriez, N., Bendidi, I., Cohen, E., Watkinson, G., Sanchez, M., Bollot, G., Genovesio, A.: Chada-vit: Channel adaptive attention for joint representation learning of heterogeneous microscopy images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11556–11565 (2024)
3. Chen, J., Ma, B., Cui, H., Xia, Y.: Think twice before selection: Federated evidential active learning for medical image analysis with domain shifts. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11439–11449 (2024)
4. Dong, J., Jiang, J., Jiang, K., Li, J., Cai, L., Zhang, Y.: Disentangled pseudo-bag augmentation for whole slide image multiple instance learning. IEEE Transactions on Medical Imaging (2025)
5. Fang, H., Huang, S., Tang, W., Huangfu, L., Liu, B.: Sam-mil: A spatial contextual aware multiple instance learning approach for whole slide image classification. In: Proceedings of the 32nd ACM international conference on multimedia. pp. 6083–6092 (2024)
6. Graham, S., Vu, Q.D., Raza, S.E.A., Azam, A., Tsang, Y.W., Kwak, J.T., Rajpoot, N.: Hover-net: Simultaneous segmentation and classification of nuclei in multi-tissue histology images. Medical image analysis **58**, 101563 (2019)
7. Huang, P., Huang, Y., Zhao, W., He, J., Yu, L.: Hyperpath: Knowledge-guided hyperbolic semantic hierarchy modeling for wsi analysis. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 262–272. Springer (2025)
8. Ilse, M., Tomczak, J., Welling, M.: Attention-based deep multiple instance learning. In: International conference on machine learning. pp. 2127–2136. PMLR (2018)
9. Jin, H., Liu, S., Cong, C., Feng, Q., Liu, Y., Huang, L., Hu, Y.: Fedwsidd: Federated whole slide image classification via dataset distillation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 178–188. Springer (2025)

10. Li, C., Liu, X., Jia, Q., Liu, J., Jiang, Z., Feng, L., Liu, Y., Luo, Z., Fan, X.: Physics-guided sonar image fine-grained recognition under scarce annotations. In: Proceedings of the 33rd ACM International Conference on Multimedia. pp. 1356–1365 (2025)
11. Li, J., Liu, J., Kuang, H., Yue, H., He, M., Wang, J.: Mico: Multiple instance learning with context-aware clustering for whole slide image analysis. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 376–385. Springer (2025)
12. Li, S., Jin, W., Wang, Z., Wu, F., Liu, Z., Tan, C., Li, S.Z.: Semireward: A general reward model for semi-supervised learning. arXiv preprint arXiv:2310.03013 (2023)
13. Li, Z., Yang, L.T., Ren, B., Nie, X., Gao, Z., Tan, C., Li, S.Z.: Mlip: Enhancing medical visual representation with divergence encoder and knowledge-guided contrastive learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11704–11714 (2024)
14. Liu, P., Ji, L., Zhang, X., Ye, F.: Pseudo-bag mixup augmentation for multiple instance learning-based whole slide image classification. IEEE Transactions on Medical Imaging **43**(5), 1841–1852 (2024)
15. Liu, Y., Li, J., Zhang, Y., Jia, Q., Wang, W., Pu, N., Sebe, N.: Pmgnet: Disentanglement and entanglement benefit mutually for compositional zero-shot learning. Computer Vision and Image Understanding **249**, 104197 (2024)
16. Lou, J., Pan, Q., Yang, Q., Ji, B.: Pcr-mil: Phenotype clustering reinforced multiple instance learning for whole slide image classification. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 395–404. Springer (2025)
17. Qu, M., Yang, G., Di, D., Su, T., Gao, Y., Song, Y., Fan, L.: Multimodal cancer survival analysis via hypergraph learning with cross-modality rebalance. arXiv preprint arXiv:2505.11997 (2025)
18. 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)
19. Sun, S., van Midden, D., Litjens, G., Baumgartner, C.F.: Prototype-based multiple instance learning for gigapixel whole slide image classification. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 507–517. Springer (2025)
20. Tian, F., Liu, D., Wei, N., Fu, Q., Sun, L., Liu, W., Sui, X., Tian, K., Nemeth, G., Feng, J., et al.: Prediction of tumor origin in cancers of unknown primary origin with cytology-based deep learning. Nature Medicine **30**(5), 1309–1319 (2024)
21. Wu, J., Chen, M., Ke, X., Xun, T., Jiang, X., Zhou, H., Shao, L., Kong, Y.: Learning heterogeneous tissues with mixture of experts for gigapixel whole slide images. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 5144–5153 (2025)
22. Wu, L., Zhuang, J., Chen, H.: Voco: A simple-yet-effective volume contrastive learning framework for 3d medical image analysis. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 22873–22882 (2024)
23. Wu, Y., Xie, X., Niu, X., Zhao, C., Yu, J.: Cross-patient pseudo bags generation and curriculum contrastive learning for imbalanced multiclassification of whole slide image. arXiv preprint arXiv:2411.11262 (2024)
24. Yan, R., Sun, Q., Jin, C., Liu, Y., He, Y., Guan, T., Chen, H.: Shapley values-enabled progressive pseudo bag augmentation for whole-slide image classification. IEEE transactions on medical imaging (2024)

25. Yang, S., Wang, Y., Chen, H.: Mambamil: Enhancing long sequence modeling with sequence reordering in computational pathology. In: International conference on medical image computing and computer-assisted intervention. pp. 296–306. Springer (2024)
26. 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)
27. Zhao, J., Li, X., Yang, F., Zhai, Q., Luo, A., Zhao, Y., Cheng, H., Fu, H.: Mexd: An expert-infused diffusion model for whole-slide image classification. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 20789–20799 (2025)
28. Zhu, C., Zhang, Y., Jia, Q., Wang, W., Liu, Y.: Temporal refinement and multi-grained matching for moment retrieval and highlight detection. Multimedia Systems **31**(1), 46 (2025)