

# Feature Re-Embedding: Towards Foundation Model-Level Performance in Computational Pathology

Wenhao Tang<sup>1†</sup> Fengtao Zhou<sup>2†</sup> Sheng Huang<sup>1\*</sup> Xiang Zhu<sup>1</sup> Yi Zhang<sup>1</sup> Bo Liu<sup>3</sup>  
<sup>1</sup> Chongqing University <sup>2</sup> The Hong Kong University of Science and Technology  
<sup>3</sup> Walmart Global Tech

{whtang, huangsheng, zhangyii}@cqu.edu.cn, zhuxiang@stu.cqu.edu.cn  
 fzhouaf@connect.ust.hk, kfliubo@gmail.com

## Abstract

Multiple instance learning (MIL) is the most widely used framework in computational pathology, encompassing sub-typing, diagnosis, prognosis, and more. However, the existing MIL paradigm typically requires an offline instance feature extractor, such as a pre-trained ResNet or a foundation model. This approach lacks the capability for feature fine-tuning within the specific downstream tasks, limiting its adaptability and performance. To address this issue, we propose a Re-embedded Regional Transformer ( $R^2T$ ) for re-embedding the instance features online, which captures fine-grained local features and establishes connections across different regions. Unlike existing works that focus on pre-training powerful feature extractor or designing sophisticated instance aggregator,  $R^2T$  is tailored to re-embed instance features online. It serves as a portable module that can seamlessly integrate into mainstream MIL models. Extensive experimental results on common computational pathology tasks validate that: 1) feature re-embedding improves the performance of MIL models based on ResNet-50 features to the level of foundation model features; 2) the  $R^2T$  can introduce more significant performance improvements to various MIL models; 3)  $R^2T$ -MIL, as an  $R^2T$ -enhanced AB-MIL, outperforms other latest methods by a large margin. The code is available at: <https://github.com/DearCaat/RRT-MIL>.

## 1. Introduction

Computational pathology [6, 7, 25] is an interdisciplinary field that combines pathology, image analysis, and computer science to develop and apply computational methods for the analysis and interpretation of pathological im-

\*Corresponding Author.

†Equal Contribution.

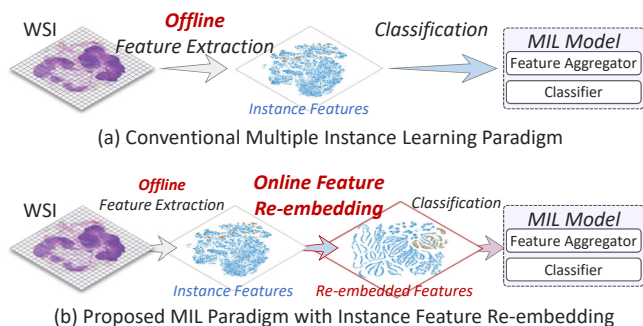


Figure 1. **Top:** The conventional MIL paradigm lacks fine-tuning of the offline embedded instance features. **Bottom:** The proposed MIL paradigm that introduces instance feature re-embedding to provide more discriminative features for the MIL model.

ages (also known as whole slide images, WSIs). This field utilizes advanced algorithms, machine learning, and artificial intelligence techniques to assist pathologists in tasks like sub-typing [12], diagnosis [16, 41], prognosis [32, 39], and more. However, the process of pixel-level labeling in ultra-high resolution WSIs is time-consuming and labor-intensive, posing challenges for traditional deep learning methods that rely on pixel-level labels in computational pathology. To address this challenge, multiple instance learning (MIL) approaches have been employed to treat WSI analysis as a weakly supervised learning problem [19, 30]. MIL divides each WSI (referred to as a bag) into numerous image patches or instances. Previous MIL-based methods mainly follow a three-step process: 1) instance feature extraction, 2) instance feature aggregation, and 3) bag prediction. However, most previous works focus on the last two steps, where the extracted offline instance features are utilized to make bag-level predictions.

Despite achieving “clinical-grade” performance on numerous computational pathology tasks [4, 41], the conventional MIL paradigm faces a significant design challenge due to the large number of instances involved. The holis-

tic end-to-end learning of the instance feature extractor, instance-level feature aggregator, and bag-level predictor becomes infeasible due to the prohibitively high memory cost. In previous works, an offline feature extractor pre-trained on natural images is used to extract instance features. However, this approach lacks a feature fine-tuning process for specific downstream tasks [18, 22, 41], resulting in low discriminative features and sub-optimal performance, as illustrated in Figure 1(a). To mitigate this issue, some works [4, 11, 14] have employed self-supervised methods to pre-train a more powerful feature extractor on a massive amount of WSIs, which are known as foundation models. Nevertheless, pre-training foundation model requires huge amounts of data (>200k WSIs) and computational resources. **Furthermore, the challenge of lacking feature fine-tuning remains unresolved.** An intuitive way of addressing the issue is to perform online features re-embedding using representation learning techniques before MIL models. As shown in Figure 1(b), re-embedding modules can be trained end-to-end with MIL models to provide supervised feature fine-tuning. It enables fully exploiting the knowledge beneficial to the final task.

As a powerful representation learning method, Transformer [29] has proven to be effective for representation learning and has demonstrated promising results in various domains [3, 38, 40]. However, directly applying the existing Transformers for re-embedding is challenging due to the characteristics of WSI. The main problem is the unacceptable memory consumption caused by the massive input of image patches. The linear multi-head self-attention (MSA) [36] can alleviate the memory dilemma, but suffers from high computational cost and sub-optimal performance. Moreover, the global MSA fails to capture the local detail features that are crucial for computational pathology.

In this paper, we propose Re-embedded Regional Transformer ( $R^2T$ ), which leverages the advantages of the native MSA while overcoming its limitations. Specifically,  $R^2T$  applies the native MSA to each local region separately. Then, it uses a Cross-region MSA (CR-MSA) to fuse the information from different regions. Finally, a novel Embedded Position Encoding Generator (EPEG) is used to effectively encode the positional information of the patches. By incorporating with mainstream MIL models, the proposed  $R^2T$  can ensure efficient computation while maintaining powerful representation capabilities to fine-tune the offline features according to the specific downstream tasks. The main contributions can be summarized as follows:

- We propose a novel paradigm for MIL models that incorporates a re-embedding module to address the issue of poor discriminative ability in instance features caused by offline feature extractors. The proposed feature re-embedding fashion can effectively improve MIL models, even achieving competitive performance compared to the

latest foundation model.

- For re-embedding instance features, we design a Re-embedded Regional Transformer ( $R^2T$ ) which can be seamlessly integrated into mainstream MIL models to further improve performance. By incorporating the  $R^2T$  into AB-MIL, we present the  $R^2T$ -MIL, which achieves state-of-the-art performance on various computational pathology benchmarks.
- We introduce two novel components for the  $R^2T$ : the Cross-region MSA and the Embedded Position Encoding Generator. The former enables effective information fusion across different regions. The latter combines the benefits of relative and convolutional position encodings to encode the positional information more effectively.

## 2. Related Work

### 2.1. Computational Pathology

The transition from traditional glass slides to digital pathology has provided a wealth of opportunities for computational pathology, which aims to combine pathology, image analysis, and computer science techniques to develop computer-assisted methods for analyzing pathology images [6, 7, 25]. By harnessing the power of advanced machine learning algorithms, computational pathology can enable large-scale data analysis and facilitate collaboration among pathologists and researchers. Traditionally, pathologists relied on visual examination of tissue samples under a microscope to make diagnoses. However, this manual process was not only time-consuming but also prone to subjective interpretations and human errors. With the emergence of computational pathology, these limitations are being addressed in remarkable ways. By automating labor-intensive processes, it can liberate pathologists' time, enabling them to focus on complex and critical decision-making tasks. Meanwhile, its ability to leverage vast amounts of data, combined with advanced analytic, holds great promise for breakthroughs in personalized medicine. By extracting quantitative features from pathology images, computational pathology can assist in making diagnosis [12, 14, 18], predicting patient outcomes [34, 46], identifying biomarkers [9, 35], and guiding tailored treatment strategies [27].

### 2.2. Multiple Instance Learning

Multiple instance learning (MIL) is the most widely used paradigm in computational pathology, involving three key steps: slide patching, instance feature extraction, and bag label prediction [2, 12, 18]. Due to the ultra-high resolution of WSIs, the instance features are typically extracted by pre-trained models, especially ResNet-50 pre-trained on ImageNet. However, the inherent difference between pathology images and nature images results in poor discrimination of extracted features. Some self-supervised

learning-based methods [4, 11, 14, 21, 43] attempt to alleviate the feature bias by pre-training feature extractor on a large number of WSIs. For example, Huang et al. adapted CLIP [20] to pre-train a vision Transformer called PLIP, with more than 200k slide-text pairs [11]. These efforts aim to enhance the discrimination of offline features by leveraging the vast amount of pathology-specific information available in the pre-training data. The extracted instance features are then utilized for bag prediction in computational pathology. These methods can be categorized into instance label fusion [2, 8, 13, 37] and instance feature fusion [14, 18, 22, 23, 41]. Instance label fusion methods first obtain instance labels and then pool them to obtain the bag label, while instance feature fusion methods aggregate all instance features into a high-level bag embedding and then obtain the bag prediction. Recently, Transformer blocks [29] have been utilized to aggregate instance features [15, 22, 31], demonstrating the advantage of self-attention over traditional attention [12, 14, 18] in modeling mutual instance information. While existing methods in computational pathology have shown promising results, most of them primarily focus on how to aggregate discriminative information from pre-extracted features. However, the pre-extracted features lack fine-tuning on specific downstream tasks, resulting in sub-optimal performance.

### 3. Methodology

#### 3.1. Preliminary

From the perspective of MIL, a WSI  $X$  is considered as a bag while its patches are deemed as instances in this bag, which can be represented as  $X = \{x_i\}_{i=1}^I$ . The instance number  $I$  varies for different bags. For a classification task, there exists a known label  $Y$  for a bag and an unknown label  $y_i$  for each of its instances. If there is at least one positive instance in a bag, then this bag is positive; otherwise, it is negative. The goal of a MIL model  $\mathcal{M}(\cdot)$  is to predict the bag label with all instances  $\hat{Y} \leftarrow \mathcal{M}(X)$ . Following the recent popular approaches [13, 37], the MIL prediction process can be divided into three steps: instance feature extraction, instance feature aggregation, and bag classification. Specifically, this process can be defined as follows:

$$\hat{Y} \leftarrow \mathcal{M}(X) := \mathcal{C}(\mathcal{A}(\mathcal{F}(X))), \quad (1)$$

where  $\mathcal{F}(\cdot)$ ,  $\mathcal{A}(\cdot)$ , and  $\mathcal{C}(\cdot)$  are the mapping functions of these aforementioned steps respectively.

In computational pathology, extracting all instances in a bag poses a huge computational challenge for the end-to-end optimization of these three steps. Therefore, most existing approaches rely on a pre-trained deep learning model to obtain instance features first. Then, they only optimize the aggregation and classification steps. **However, the non-fine-tuned features lead to sub-optimal performance,**

**even if the features are extracted by a foundation model.**

An intuitive way to address this problem is by re-embedding based on the extracted instance features, while most of the existing approaches pay more attention to feature aggregation and neglect the importance of re-embedding. In this paper, we include a re-embedding step after the instance feature extraction and update the bag labeling as follows,

$$\hat{Y} \leftarrow \mathcal{M}(X) := \mathcal{C}(\mathcal{A}(\mathcal{R}(\mathcal{F}(X)))), \quad (2)$$

where  $\mathcal{R}(\cdot)$  is the mapping function of the re-embedding.

#### 3.2. Re-embedded Regional Transformer

As illustrated in Figure 2, we propose a Re-embedded Regional Transformer (R<sup>2</sup>Transformer, R<sup>2</sup>T) to re-embed the input instance features as the new instance representations,

$$Z = \{z_i\}_{i=1}^I := \mathcal{R}(H) \in \mathbb{R}^{I \times D}, \quad (3)$$

where  $\mathcal{R}(\cdot)$  is the mapping function of the R<sup>2</sup>Transformer here,  $H = \{h_i\}_{i=1}^I := \mathcal{F}(X) \in \mathbb{R}^{I \times D}$  is the processed input instance features, and  $z_i = \mathcal{R}(h_i)$  is the new embedding of the  $i$ -th instance.  $D$  is the dimension of the embedding. The R<sup>2</sup>Transformer can be flexibly plugged into the MIL framework as a re-embedding module after feature input and before instance aggregation to reduce bias caused by the shift between offline feature learning and downstream tasks. The whole re-embedding process in the R<sup>2</sup>Transformer can be formulated as,

$$\begin{aligned} \hat{Z} &= \text{R-MSA}(\text{LN}(H)) + H \\ Z &= \text{CR-MSA}\left(\text{LN}\left(\hat{Z}\right)\right) + \hat{Z} \end{aligned} \quad (4)$$

where R-MSA( $\cdot$ ) denotes Regional Multi-head Self-attention, CR-MSA( $\cdot$ ) denotes Cross-region MSA, and LN( $\cdot$ ) denotes Layer Normalization.

**Regional Multi-head Self-attention:** Since instance number  $I$  is very large, most Transformers in this field commonly adopt two strategies to avoid the Out-of-Memory issue. The first method is to sample or aggregate the original large instance set into a small one, after which global self-attention is performed [15, 44]. The second method performs the Nystrom algorithm [36] to approximate the global self-attention [22]. Although these methods address the scalability issue of self-attention with a large  $I$ , they neglect the fact that the tumor areas are local and only occupy a small part of the whole image. Performing global self-attention on all instances results in feature homogenization. Moreover, different from the conventional MIL application scenarios, the instances in each bag have ordinal relations in computational pathology due to the fact that they are all collected from the same slide in order. These facts motivate us to design the Regional Multi-head Self-attention (R-MSA)




Figure 2. Overview of proposed  $R^2T$ -MIL. A set of patches is first cropped from the tissue regions of a slide and embedded in features by an offline extractor. Then, the sequence is processed with the  $R^2T$  module: (1) region partition, (2) feature re-embedding within each region, and (3) cross-region feature fusion. Finally, a MIL model predicts the bag labels using the re-embedded instance features.

that divides the bag into several different regions and performs self-attention in each region separately. R-MSA takes into account the aforementioned WSI properties and makes use of instance ordinal relation information to reduce computation complexity and highlight salient local features.

In R-MSA, the input instance features are reshaped into a 2-D feature map,  $H \in \mathbb{R}^{I \times D} \rightarrow H \in \mathbb{R}^{\lceil \sqrt{I} \rceil \times \lceil \sqrt{I} \rceil \times D}$ . And  $L \times L$  regions are then divided evenly across the map in a non-overlapping manner, with each containing  $M \times M$  instances where  $L \times M = \lceil \sqrt{I} \rceil$ . For example, the region partition starts from the top-left instance, and an  $8 \times 8$  feature map is evenly partitioned into  $2 \times 2$  regions of size  $4 \times 4$  ( $L = 2, M = 4$ ). We fix the number of regions  $L$  rather than the size of regions  $M$  to obtain  $L \times L$  regions with adaptive size. By default,  $L$  is set to 8. Self-attention is computed within each local region. The whole process of R-MSA can be denoted as,

$$\begin{aligned}
 \text{Step 1: } & H \in \mathbb{R}^{I \times D} \xrightarrow{\text{Squaring}} H \in \mathbb{R}^{L^2 \times M^2 \times D}, \\
 \text{Step 2: } & H \xrightarrow{\text{Partition}} \{H^l\}_{l=1}^{L^2}, H^l \in \mathbb{R}^{M \times M \times D}, \\
 \text{Step 3: } & \hat{Z} := \{\hat{Z}^l\}_{l=1}^{L^2}, \hat{Z}^l = S(H^l) \in \mathbb{R}^{M \times M \times D},
 \end{aligned} \quad (5)$$

where  $S(\cdot)$  is vanilla multi-head self-attention with our proposed Embedded Position Encoding Generator (EPEG).

**Embedded Position Encoding Generator:** Inspired by [22], we adopt a convolutional computation called Position Encoding Generator (PEG) [5] to address the challenge of traditional position encoding strategies being unable to handle input sequences of variable length. Different from previous methods, we propose a novel approach called Embedded PEG (EPEG) by incorporating PEG into the R-MSA module, inspired by the relative position encoding strategy [17, 24, 33]. By embedding the PEG into the MSA module, EPEG can utilize a lightweight 1-D convolution  $\text{Conv}_{1-D}(\cdot)$  to more effectively encode in each region

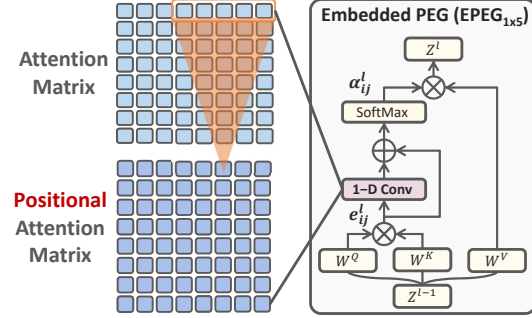


Figure 3. Illustration of Embedded Position Encoding Generator.

separately. The structure of EPEG is shown in Figure 3. Taking the instances in the  $l$ -th region as an example, the EPEG can be formulated as,

$$\alpha_{ij}^l = \text{SoftMax} (e_{ij}^l + \text{Conv}_{1-D} (e_{ij}^l)), \quad (6)$$

where  $\alpha_{ij}^l$  is the attention weight of  $H_j^l$  with respect to  $H_i^l$ , and  $e_{ij}^l$  is calculated using a scaled dot-product attention.

**Cross-region Multi-head Self-attention:** R-MSA only considers the features within each region, which limits its modeling power of context-based semantic features. This is crucial for downstream tasks, such as prognosis, which requires a more comprehensive judgment. To effectively model the cross-region connections, we propose Cross-region Multi-head Self-attention (CR-MSA). First, we aggregate the representative features  $R^l$  of each region,

$$\begin{aligned}
 W_a^l &= \text{SoftMax}_{m=1}^M \left( \hat{Z}_m^l \Phi \right), \\
 R^l &= W_a^{lT} \hat{Z}^l,
 \end{aligned} \quad (7)$$

where  $\Phi \in \mathbb{R}^{D \times K}$  denotes learnable parameters. We utilize vanilla MSA to model the cross-region connection,  $\hat{R} = S(R)$ . Finally, the updated representative features

Methods		CAMELYON-16			TCGA-BRCA		
		Accuracy	AUC	F1-score	Accuracy	AUC	F1-score
ResNet-50 ImageNet pretrained	AB-MIL [12]	90.06±0.72	94.54±0.30	87.83±0.83	86.41±4.92	91.10±2.52	81.64±4.71
	CLAM [18]	90.14±0.85	94.70±0.76	88.10±0.63	85.17±2.70	91.67±1.78	80.37±3.04
	DSMIL [14]	90.17±1.02	94.57±0.40	87.65±1.18	<u>87.20±2.69</u>	91.58±1.33	82.41±2.92
	TransMIL [22]	89.22±2.32	93.51±2.13	85.10±4.33	<u>84.68±2.67</u>	90.80±1.91	79.86±2.63
	DTFD-MIL [41]	90.22±0.36	95.15±0.14	87.62±0.59	85.92±1.76	91.43±1.64	81.09±2.05
	IBMIL [16]	91.23±0.41	94.80±1.03	88.80±0.89	84.19±3.40	91.01±2.32	79.45±3.42
	MHIM-MIL [26]	<u>91.81±0.82</u>	<u>96.14±0.52</u>	<u>89.94±0.70</u>	86.73±5.59	<u>92.36±1.58</u>	<u>82.43±5.47</u>
<b>R<sup>2</sup>T-MIL</b>	<b>92.40±0.31</b>	<b>97.32±0.29</b>	<b>90.63±0.45</b>	<b>88.33±0.67</b>	<b>93.17±1.45</b>	<b>83.70±0.95</b>	
PLIP WSI pretrained	AB-MIL [12]	94.66±0.42	97.30±0.31	93.29±0.54	85.45±2.32	91.73±2.26	80.60±2.66
	CLAM [18]	93.73±0.54	97.17±0.50	91.60±0.60	86.70±1.35	92.16±2.02	81.91±1.78
	DSMIL [14]	94.40±0.85	97.06±0.56	92.78±1.15	87.25±2.70	91.80±1.67	82.18±2.28
	TransMIL [22]	94.40±0.43	<u>97.88±0.21</u>	92.81±0.43	85.83±3.44	92.17±2.20	81.12±3.25
	DTFD-MIL [41]	94.57±0.31	97.29±0.23	93.12±0.40	86.42±2.67	92.16±2.42	81.77±2.73
	IBMIL [16]	93.90±0.66	97.04±0.18	92.44±0.64	<u>87.57±1.48</u>	91.71±1.74	<u>82.78±2.02</u>
	MHIM-MIL [26]	<u>95.32±0.31</u>	97.79±0.15	94.13±0.42	<u>87.07±2.20</u>	93.17±2.00	82.48±2.50
<b>R<sup>2</sup>T-MIL</b>	<b>95.49±0.00</b>	<b>98.05±0.29</b>	<b>94.29±0.04</b>	<b>88.82±3.22</b>	<b>93.80±1.24</b>	<b>84.55±3.55</b>	

Table 1. Cancer diagnosis, and sub-typing results on C16 and BRCA. The highest performance is in bold, and the second-best performance is underlined. With AB-MIL as a baseline, R<sup>2</sup>T-MIL is not only capable of re-embedding ResNet-50 features to the level of foundation model (PLIP [11]) features, but also effectively fine-tuning offline PLIP features.

are distributed to each instance in the region with MinMax normalized weight  $W_d^l$ ,

$$W_d^l = \text{MinMax}_{m=1}^M \left( \hat{Z}_m^l \Phi \right) \in \mathbb{R}^{M^2 \times K}, \quad (8)$$

$$Z^l = W_d^{l\top} \hat{R}^l \hat{W}_d^l,$$

where  $\hat{W}_d^l = \text{SoftMax}_{k=1}^K \left( \hat{Z}_{mk}^l \Phi \right) \in \mathbb{R}^{K \times 1}$ .

### 3.3. R<sup>2</sup>Transformer-based MIL

Once we obtain the re-embedding of instances, any instance aggregation method and classifier can be applied to accomplish the specific downstream tasks. The re-embedding  $\mathcal{R}(\cdot)$  will be optimized with the instance aggregation module  $\mathcal{A}(\cdot)$  and the bag classifier  $\mathcal{C}(\cdot)$  together,

$$\{\hat{\mathcal{R}}, \hat{\mathcal{A}}, \hat{\mathcal{C}}\} \leftarrow \arg \min L(Y, \hat{Y}) = L(Y, \mathcal{C}(\mathcal{A}(\mathcal{R}(H))))), \quad (9)$$

where  $L(\cdot, \cdot)$  denotes any MIL loss. R<sup>2</sup>Transformer-based MIL (R<sup>2</sup>T-MIL) adopts the instance aggregation method and the bag classifier of AB-MIL [12] by default.

## 4. Experiments and Results

### 4.1. Datasets and Evaluation Metrics

**Datasets:** We use **CAMELYON-16** [1] (C16), **TCGA-BRCA**, and **TCGA-NSCLC** to evaluate the performance on diagnosis and sub-typing tasks. For prognosis, we use **TCGA-LUAD**, **TCGA-LUSC**, **TCGA-BLCA** to evaluate the performance on the survival prediction task. Please refer to the Supplementary Material for more details.

**Evaluation Metrics:** For diagnosis and sub-typing, we leverage Accuracy, Area Under Curve (AUC), and F1-score to evaluate model performance. We only report AUC in ablation experiments. For survival prediction, we report the C-index in all datasets. To reduce the impact of data split on model evaluation, we follow [18, 42, 45] and apply 5-fold cross-validation in all remaining datasets except C16. We report the mean and standard deviation of the metrics over  $N$  folds. For C16, we follow [26] and use 3-times 3-fold cross-validation to alleviate the effects of random seed.

**Compared Methods:** Seven influential MIL approaches are employed for comparison. They are AB-MIL [12], DSMIL [14], CLAM [18], DTFD-MIL [41], TransMIL [22], IBMIL [16], and MHIM-MIL [26]. We reproduce the results of these methods under the same settings.

**Implementation Details:** We adopt ResNet50 [10] pre-trained with ImageNet-1k and the latest foundation model PLIP [11] pre-trained with OpenPath as the offline feature extractors. Supplementary Material offers more details.

## 4.2. Main Results

### 4.2.1 Cancer Diagnosis, and Sub-typing

Table 1 presents the diagnosis and sub-typing performances of different MIL approaches on the C16 and BRCA datasets. The results demonstrate that our proposed R<sup>2</sup>T-MIL achieves the best performance under all metrics on all benchmarks. Specifically, R<sup>2</sup>T-MIL gets 0.59%, 1.18%, and 0.69% performance gains over the second-best methods in Accuracy, AUC, and F1-score respectively on the

	Methods	Accuracy	AUC	F1-score
ResNet-50	AB-MIL	90.32 $\pm$ 1.39	95.29 $\pm$ 1.14	89.83 $\pm$ 1.53
	CLAM	90.52 $\pm$ 2.08	95.37 $\pm$ 1.08	90.08 $\pm$ 1.97
	DSMIL	90.43 $\pm$ 2.52	95.60 $\pm$ 0.81	90.03 $\pm$ 2.61
	TransMIL	90.04 $\pm$ 1.86	94.97 $\pm$ 1.11	89.94 $\pm$ 1.73
	DTFD-MIL	89.85 $\pm$ 1.53	95.55 $\pm$ 1.47	89.60 $\pm$ 1.67
	IBMIL	90.04 $\pm$ 1.48	95.57 $\pm$ 1.13	89.73 $\pm$ 1.64
	MHIM-MIL	91.27 $\pm$ 2.35	96.02 $\pm$ 1.35	90.85 $\pm$ 2.53
	<b>R<sup>2</sup>T-MIL</b>	<b>91.75<math>\pm</math>2.38</b>	<b>96.40<math>\pm</math>1.13</b>	<b>91.26<math>\pm</math>2.60</b>
PLIP	AB-MIL	90.99 $\pm$ 2.43	95.68 $\pm$ 1.98	90.52 $\pm$ 2.45
	CLAM	90.80 $\pm$ 2.35	95.46 $\pm$ 1.72	90.38 $\pm$ 2.46
	DSMIL	90.60 $\pm$ 2.37	95.78 $\pm$ 1.81	90.24 $\pm$ 2.51
	TransMIL	89.09 $\pm$ 2.00	95.30 $\pm$ 1.95	88.83 $\pm$ 2.16
	DTFD-MIL	90.42 $\pm$ 2.98	95.83 $\pm$ 1.75	89.91 $\pm$ 3.01
	IBMIL	91.18 $\pm$ 3.27	95.62 $\pm$ 2.09	90.94 $\pm$ 3.20
	MHIM-MIL	91.74 $\pm$ 1.88	96.21 $\pm$ 1.26	91.20 $\pm$ 1.89
	<b>R<sup>2</sup>T-MIL</b>	<b>92.13<math>\pm</math>2.55</b>	<b>96.40<math>\pm</math>1.45</b>	<b>91.83<math>\pm</math>2.50</b>

Table 2. Sub-typing results on TCGA-NSCLC.

C16 dataset. On the BRCA dataset, the AUC improvement is 0.69%. R<sup>2</sup>T-MIL employs the same aggregation and classification methods as AB-MIL. However, R<sup>2</sup>T-MIL significantly outperforms AB-MIL. It increases the AUC by 2.78% and 1.77% on the C16 and BRCA datasets, respectively. The sub-typing results on NSCLC in Table 2 support a similar observation. We attribute these substantial performance improvements to the additional re-embedding step based on our proposed R<sup>2</sup>T, which surpasses the performance of the foundation model (+0.02% AUC on C16, +1.37% on BRCA, +0.72% on NSCLC). In addition, we find that R<sup>2</sup>T can further enhance the features of the foundation model, achieving considerable improvement. This validates the effectiveness of re-embedding.

#### 4.2.2 Survival Prediction

Table 3 shows the experimental results on three survival prediction datasets. It is worth noting that our proposed R<sup>2</sup>T-MIL model demonstrates outstanding performance, attaining a C-index of 61.13% on the BLCA, 67.19% on the LUAD, and 60.95% on the LUSC. It outperforms the compared methods by a significant margin, with improvements of 2.23%, 3.08%, and 1.77% over the second-best methods, respectively. Furthermore, our proposed feature re-embedding strategy can yield substantial improvements even when working with high-quality features extracted by the foundation model. Particularly, compared to AB-MIL, our feature re-embedding strategy brings performance improvements of 4.8%, 3.85%, and 3.3% for the three datasets, respectively. These results highlight the consistent and reliable performance of our proposed strategy and method, indicating its efficacy in predicting survival outcomes.

	Methods	BLCA	LUAD	LUSC
ResNet-50	AB-MIL	57.50 $\pm$ 3.94	58.78 $\pm$ 4.90	56.51 $\pm$ 7.14
	CLAM	57.57 $\pm$ 3.73	59.60 $\pm$ 3.93	56.65 $\pm$ 6.90
	DSMIL	57.42 $\pm$ 2.25	59.31 $\pm$ 4.75	55.03 $\pm$ 6.61
	TransMIL	58.90 $\pm$ 4.70	64.11 $\pm$ 1.99	56.39 $\pm$ 2.94
	DTFD-MIL	56.98 $\pm$ 3.24	59.48 $\pm$ 2.61	55.16 $\pm$ 4.33
	IBMIL	58.41 $\pm$ 2.90	58.58 $\pm$ 4.67	59.18 $\pm$ 3.29
	MHIM-MIL	58.36 $\pm$ 3.26	60.32 $\pm$ 4.41	56.08 $\pm$ 6.33
	<b>R<sup>2</sup>T-MIL</b>	<b>61.13<math>\pm</math>2.36</b>	<b>67.19<math>\pm</math>4.02</b>	<b>60.95<math>\pm</math>4.41</b>
PLIP	AB-MIL	59.18 $\pm$ 2.48	62.09 $\pm$ 4.38	57.12 $\pm$ 2.39
	CLAM	61.58 $\pm$ 2.89	64.05 $\pm$ 4.70	58.00 $\pm$ 3.34
	DSMIL	58.96 $\pm$ 1.80	63.82 $\pm$ 5.56	56.12 $\pm$ 2.21
	TransMIL	56.20 $\pm$ 3.26	63.55 $\pm$ 2.94	58.84 $\pm$ 3.28
	DTFD-MIL	59.67 $\pm$ 4.71	61.78 $\pm$ 2.33	57.75 $\pm$ 3.52
	IBMIL	56.32 $\pm$ 2.69	58.86 $\pm$ 3.40	57.33 $\pm$ 3.28
	MHIM-MIL	60.92 $\pm$ 3.38	62.94 $\pm$ 4.58	55.95 $\pm$ 2.54
	<b>R<sup>2</sup>T-MIL</b>	<b>63.98<math>\pm</math>2.26</b>	<b>65.94<math>\pm</math>1.34</b>	<b>60.42<math>\pm</math>2.15</b>

Table 3. Survival Prediction results on three main datasets.

### 4.3. Ablation Study

#### 4.3.1 Re-Embedding Matters in Pathology

**Foundation Model Features vs. Re-embedding Features:** Table 4 and Figure 4 compare the performance of the features extracted by the foundation model PLIP [11] and various re-embedding features. The PLIP is based on the multi-modal foundation model CLIP [20] and uses up to 200K slide-text pairs for pre-training. Although this high-cost pre-training brings some improvement on different tasks, it still has bottlenecks. We attribute this to the dilemma of the traditional paradigm that even the best offline pre-training features cannot address the issue of insufficient feature fine-tuning for downstream tasks. In contrast, re-embedding modules that can be end-to-end trained with MIL models provide supervised feature fine-tuning, which enables full exploitation of the knowledge beneficial to the final task. Hence, we can see from Table 4 that any re-embedding structure achieves a considerable improvement on all tasks. Some tailored structures, such as our proposed R<sup>2</sup>T, can have a significant performance advantage in most of the tasks. Moreover, Figure 4 shows that this improvement is not limited to the classical AB-MIL, but can widely benefit different MIL models. Table 1 demonstrates that the re-embedding approach is still effective on foundation model features. Therefore, compared to foundation model features, re-embedding features are a cheaper, more versatile alternative and effective booster.

**Different Re-embedding Discussion:** The bottom part of Table 4 presents the performance of AB-MIL under different feature settings. We employ four methods, including TranMIL, N-MSA, the local version of N-MSA, and our proposed R<sup>2</sup>T to re-embed the features. From observa-




Figure 4. Performance improvement by adding  $R^2T$ . Features re-embedded by  $R^2T$  online outperform PLIP offline features on most tasks.

Model	C16 $\uparrow$	NSCLC $\uparrow$	LUAD $\uparrow$	$TT_{C16}$ $\downarrow$
AB-MIL+R50	94.52	95.28	58.78	3.1s
AB-MIL+PLIP	97.30	95.68	62.09	-
<i>R50+Re-embedding</i>				
+TransMIL (global)	95.80	95.58	63.24	13.2s
+N-MSA (global)	96.20	95.51	63.99	7.7s
+N-MSA (local)	96.47	95.97	65.41	29.8s
<b>+<math>R^2T</math> (local)</b>	<b>97.32</b>	<b>96.40</b>	<b>67.19</b>	<b>6.5s</b>

Table 4. Comparison of different instance features under AB-MIL. We report the train time per epoch on C16 ( $TT_{C16}$ ). The pre-training time is not included for comparison.

tions, all four re-embedded AB-MILs perform better than the original one. The performance improvement by TransMIL, N-MSA, and  $R^2T$  are 1.28%, 1.68%, and 2.80%, respectively, on C16, while these numbers on LUAD are 4.46%, 4.14%, and 8.41%, respectively. This phenomenon validates the importance of re-embedding in MIL-based computational pathology. Among the four employed re-embedding approaches,  $R^2T$  boosts the AB-MIL with the most considerable improvements while incurring the lowest computational cost. Specifically,  $R^2T$  achieves more performance gains (+0.85% on C16 and +2.63% on LUAD) while only requiring a 1/5 inference time compared with the second-best approaches.

**The Applicability of  $R^2T$  in MIL Frameworks:** We incorporate the  $R^2T$  into different MIL frameworks as a re-embedding module for studying its applicability. The performance improvements achieved by this re-embedding module in different frameworks are shown in Figure 4. The results reveal that the  $R^2T$  is capable of improving all MIL frameworks on all tasks. Moreover, the improvement brought by  $R^2T$  surpasses the foundation model features on two tasks except for diagnosis, and reaches a similar level on C16. This clearly verifies the good applicability of  $R^2T$ .

#### 4.3.2 Local is more Appropriate than Global

To investigate the role of local self-attention in computational pathology, we replace the naive MSA in the partitioned regions with N-MSA [36]. The results in Table 4 demonstrate that feature re-embedding within local regions,

Model	C16 $\uparrow$	NSCLC $\uparrow$	LUAD $\uparrow$
w/o	96.82	96.01	65.45
PEG $_{3\times 3}$	96.86 (+0.04)	96.11 (+0.10)	65.61 (+0.16)
PEG $_{7\times 7}$	95.47 (-1.35)	95.94 (-0.07)	65.14 (-0.31)
PPEG	93.00 (-3.82)	96.03 (+0.02)	65.28 (-0.17)
<b>EPEG</b>	<b>97.32 (+0.50)</b>	<b>96.40 (+0.39)</b>	<b>67.19 (+1.74)</b>

Table 5. Comparison results of different position encoding. The PPEG [22] consists of a  $3 \times 3$ ,  $5 \times 5$ , and  $7 \times 7$  convolution block.

called +N-MSA (local), outperforms the global method under the same N-MSA. This validates the superiority of local self-attention in mining fine-grained features over global ones. The cost of the performance improvement in local self-attention brings a new problem. The local ones suffer from a higher computational burden than the global ones (around  $4\times$  more training time). Our proposed  $R^2T$  can alleviate this problem since it employs a more naive MSA, which significantly reduces the computational cost. Another advantage of local self-attention is that it improves the diversity of the re-embedded features, as shown in Figure 5 (c) and (d). Throughout, the local self-attention fashion is far more appropriate for re-embedding in computational pathology, and our proposed  $R^2T$  ensures both good performance and good efficiency in local self-attention.

#### 4.3.3 Effects of EPEG

We discuss the impact of various positional encoding methods that can handle variable input lengths in detail here. Table 5 shows that the conventional conditional positional encoding methods for the input, such as PEG [5] and PPEG [22], do not effectively improve the performance. More parameters and more complex structures do not bring significant improvements, but rather the simplest PEG $_{3\times 3}$  achieves slight improvements on both tasks. In contrast, the EPEG, which is embedded in MSA, can benefit from a more lightweight 1-D convolution, and encode the attention matrix more directly. This enables it to model the positional information more effectively in the re-embedding module. For instance, EPEG obtains 0.50%, 0.39%, and 1.74% im-




Figure 5. The tSNE [28] visualization of instance features from the CAMELYON-16 dataset, comparing (a) features extracted by ResNet-50 pre-trained on ImageNet-1k, (b) features extracted by PLIP, (c) features after N-MSA re-embedding, and (d) features after R<sup>2</sup>T re-embedding. In (a), we obtain instance-level labels from the tumor annotations and report the instance numbers of different labels.


provements on C16, NSCLC, and LUAD, respectively.

#### 4.3.4 Impact of Cross-Region MSA

Although R-MSA can effectively mine the fine-grained features of local regions, the hard partitioning would restrict the range of re-embedding to each separate region. This impairs the discriminative power of the features, as they lack cross-region connections. The left figure in Table 6 illustrates this phenomenon, where all features are divided into 64 clusters (corresponding to the number of regions). Even though one cluster can capture fine-grained features, the key clusters are scattered and not cohesive. This affects the expression of context-based semantic information, which is crucial for downstream tasks. The right figure shows the significant improvement of the feature distribution after adding the CR-MSA module, and the table results also prove its effectiveness. Moreover, such multi-region-based semantic information is more important for survival prediction than the other two tasks. This is because this task is based on cases, where each bag contains multiple slides, which requires a more comprehensive discrimination.

#### 4.4. Visualization

We visualize the features before and after re-embedding with different ways in Figure 5. From observations, we can summarize that: 1) The offline feature extractor fails to learn the discriminative instance features, even with a foundation model trained on 200K slide-text pairs, especially when the instance distribution is extremely imbalanced, e.g., 1:224 positive-to-negative ratio; 2) Global self-attention enables the re-embedding of instance features, but its attention distribution is almost uniform. This indicates that its re-embedded features are homogeneous and lack diversity, which limits the performance of advanced MIL



Model	C16↑	NSCLC↑	LUAD↑
w/o	96.89	96.24	63.03
<b>w/ CR-MSA</b>	<b>97.32 (+0.43)</b>	<b>96.40 (+0.16)</b>	<b>67.19 (+4.16)</b>

Table 6. Quantitative and qualitative analysis of CR-MSA.

models. 3) The features re-embedded by R<sup>2</sup>T not only enhance discriminability but also alleviate the issue of feature homogenization.

## 5. Conclusion

In this work, we have demonstrated the importance of instance feature re-embedding for computational pathology algorithms based on MIL, alleviating the issue of the under-learning of instance features in the conventional MIL paradigm. We have also shown that Transformer-based re-embedding modules can consistently boost the performance of various MIL methods regardless of their architectures. However, the main result of this paper is the introduction of the Re-embedded Regional Transformer and two novel components: CR-MSA and EPEG. We have evidence of the importance of the local Transformer in the age of the foundation model and its versatility as a re-embedding module.



## References

- [1] Babak Ehteshami Bejnordi, Mitko Veta, Paul Johannes Van Diest, Bram Van Ginneken, Nico Karssemeijer, Geert Litjens, Jeroen AWM Van Der Laak, Meyke Hermsen, Quirine F Manson, Maschenka Balkenhol, et al. Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer. *JAMA*, 318(22):2199–2210, 2017. [5](#)
- [2] Gabriele Campanella, Matthew G Hanna, Luke Geneslaw, Allen Mirafior, Vitor Werneck Krauss Silva, Klaus J Busam, Edi Brogi, Victor E Reuter, David S Klimstra, and Thomas J Fuchs. Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nature Medicine*, 25(8):1301–1309, 2019. [2](#), [3](#)
- [3] Dexiong Chen, Leslie O’Bray, and Karsten Borgwardt. Structure-aware transformer for graph representation learning. In *International Conference on Machine Learning*, pages 3469–3489. PMLR, 2022. [2](#)
- [4] Richard J Chen, Chengkuan Chen, Yicong Li, Tiffany Y Chen, Andrew D Trister, Rahul G Krishnan, and Faisal Mahmood. Scaling vision transformers to gigapixel images via hierarchical self-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16144–16155, 2022. [1](#), [2](#), [3](#)
- [5] Xiangxiang Chu, Zhi Tian, Bo Zhang, Xinlong Wang, Xiaolin Wei, Huaxia Xia, and Chunhua Shen. Conditional positional encodings for vision transformers. *arXiv preprint arXiv:2102.10882*, 2021. [4](#), [7](#)
- [6] Didem Cifci, Gregory P Veldhuizen, Sebastian Foersch, and Jakob Nikolas Kather. Ai in computational pathology of cancer: Improving diagnostic workflows and clinical outcomes? *Annual Review of Cancer Biology*, 7:57–71, 2023. [1](#), [2](#)
- [7] Miao Cui and David Y Zhang. Artificial intelligence and computational pathology. *Laboratory Investigation*, 101(4):412–422, 2021. [1](#), [2](#)
- [8] Ji Feng and Zhi-Hua Zhou. Deep miml network. In *Proceedings of the AAAI conference on artificial intelligence*, 2017. [3](#)
- [9] Peter Hamilton, Paul O’Reilly, Peter Bankhead, Esther Abels, and Manuel Salto-Tellez. Digital and computational pathology for biomarker discovery. *Predictive Biomarkers in Oncology: Applications in Precision Medicine*, pages 87–105, 2019. [2](#)
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. [5](#)
- [11] Zhi Huang, Federico Bianchi, Mert Yuksekogul, Thomas J Montine, and James Zou. A visual–language foundation model for pathology image analysis using medical twitter. *Nature Medicine*, pages 1–10, 2023. [2](#), [3](#), [5](#), [6](#)
- [12] Maximilian Ilse, Jakub Tomczak, and Max Welling. Attention-based deep multiple instance learning. In *ICML*, pages 2127–2136. PMLR, 2018. [1](#), [2](#), [3](#), [5](#)
- [13] Fahdi Kanavati, Gouji Toyokawa, Seiya Momosaki, Michael Rambeau, Yuka Kozuma, Fumihiko Shoji, Koji Yamazaki, Sadanori Takeo, Osamu Iizuka, and Masayuki Tsuneki. Weakly-supervised learning for lung carcinoma classification using deep learning. *Scientific reports*, 10(1):9297, 2020. [3](#)
- [14] Bin Li, Yin Li, and Kevin W Eliceiri. Dual-stream multiple instance learning network for whole slide image classification with self-supervised contrastive learning. In *CVPR*, pages 14318–14328, 2021. [2](#), [3](#), [5](#)
- [15] Hang Li, Fan Yang, Yu Zhao, Xiaohan Xing, Jun Zhang, Mingxuan Gao, Junzhou Huang, Liansheng Wang, and Jianhua Yao. Dt-mil: Deformable transformer for multi-instance learning on histopathological image. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 206–216. Springer, 2021. [3](#)
- [16] Tiancheng Lin, Zhimiao Yu, Hongyu Hu, Yi Xu, and Changwen Chen. Interventional bag multi-instance learning on whole-slide pathological images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19830–19839, 2023. [1](#), [5](#)
- [17] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021. [4](#)
- [18] Ming Y Lu, Drew FK Williamson, Tiffany Y Chen, Richard J Chen, Matteo Barbieri, and Faisal Mahmood. Data-efficient and weakly supervised computational pathology on whole-slide images. *Nature Biomedical Engineering*, 5(6):555–570, 2021. [2](#), [3](#), [5](#)
- [19] Ming Y Lu, Drew FK Williamson, Tiffany Y Chen, Richard J Chen, Matteo Barbieri, and Faisal Mahmood. Data-efficient and weakly supervised computational pathology on whole-slide images. *Nature biomedical engineering*, 5(6):555–570, 2021. [1](#)
- [20] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. [3](#), [6](#)
- [21] Charlie Saillard, Olivier Dehaene, Tanguy Marchand, Olivier Moindrot, Aurélie Kamoun, Benoit Schmauch, and Simon Jegou. Self supervised learning improves dmmr/msi detection from histology slides across multiple cancers. *arXiv preprint arXiv:2109.05819*, 2021. [3](#)
- [22] Zhuchen Shao, Hao Bian, Yang Chen, Yifeng Wang, Jian Zhang, Xiangyang Ji, et al. Transmil: Transformer based correlated multiple instance learning for whole slide image classification. *NeurIPS*, 34, 2021. [2](#), [3](#), [4](#), [5](#), [7](#)
- [23] Yash Sharma, Aman Shrivastava, Lubaina Ehsan, Christopher A Moskaluk, Sana Syed, and Donald E Brown. Cluster-to-conquer: A framework for end-to-end multi-instance learning for whole slide image classification. *arXiv preprint arXiv:2103.10626*, 2021. [3](#)
- [24] Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. Self-attention with relative position representations. *arXiv preprint arXiv:1803.02155*, 2018. [4](#)

- [25] Andrew H Song, Guillaume Jaume, Drew FK Williamson, Ming Y Lu, Anurag Vaidya, Tiffany R Miller, and Faisal Mahmood. Artificial intelligence for digital and computational pathology. *Nature Reviews Bioengineering*, pages 1–20, 2023. 1, 2
- [26] Wenhao Tang, Sheng Huang, Xiaoxian Zhang, Fengtao Zhou, Yi Zhang, and Bo Liu. Multiple instance learning framework with masked hard instance mining for whole slide image classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4078–4087, 2023. 5
- [27] William T Tran, Katarzyna Jerzak, Fang-I Lu, Jonathan Klein, Sami Tabbarah, Andrew Lagree, Tina Wu, Ivan Rosado-Mendez, Ethan Law, Khadijeh Saednia, et al. Personalized breast cancer treatments using artificial intelligence in radiomics and pathomics. *Journal of medical imaging and radiation sciences*, 50(4):S32–S41, 2019. 2
- [28] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008. 8
- [29] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 2, 3
- [30] Xi Wang, Hao Chen, Caixia Gan, Huangjing Lin, Qi Dou, Efstratios Tsougenis, Qitao Huang, Muyan Cai, and Pheng-Ann Heng. Weakly supervised deep learning for whole slide lung cancer image analysis. *IEEE transactions on cybernetics*, 50(9):3950–3962, 2019. 1
- [31] Zhihua Wang, Lequan Yu, Xin Ding, Xuehong Liao, and Liansheng Wang. Lymph node metastasis prediction from whole slide images with transformer-guided multi-instance learning and knowledge transfer. *IEEE Transactions on Medical Imaging*, 2022. 3
- [32] Zhuoyu Wen, Shidan Wang, Donghan M Yang, Yang Xie, Mingyi Chen, Justin Bishop, and Guanghua Xiao. Deep learning in digital pathology for personalized treatment plans of cancer patients. In *Seminars in Diagnostic Pathology*, pages 109–119. Elsevier, 2023. 1
- [33] Kan Wu, Houwen Peng, Minghao Chen, Jianlong Fu, and Hongyang Chao. Rethinking and improving relative position encoding for vision transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10033–10041, 2021. 4
- [34] Ellery Wulczyn, David F Steiner, Zhaoyang Xu, Apaar Sathwani, Hongwu Wang, Isabelle Flament-Auvigne, Craig H Mermel, Po-Hsuan Cameron Chen, Yun Liu, and Martin C Stumpe. Deep learning-based survival prediction for multiple cancer types using histopathology images. *PloS one*, 15(6):e0233678, 2020. 2
- [35] Xiaoliang Xie, Xulin Wang, Yuebin Liang, Jingya Yang, Yan Wu, Li Li, Xin Sun, Pingping Bing, Binsheng He, Geng Tian, et al. Evaluating cancer-related biomarkers based on pathological images: a systematic review. *Frontiers in Oncology*, 11:763527, 2021. 2
- [36] Yunyang Xiong, Zhanpeng Zeng, Rudrasis Chakraborty, Mingxing Tan, Glenn Fung, Yin Li, and Vikas Singh. Nyströmformer: A nyström-based algorithm for approximating self-attention. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 14138–14148, 2021. 2, 3, 7
- [37] Gang Xu, Zhigang Song, Zhuo Sun, Calvin Ku, Zhe Yang, Cancheng Liu, Shuhao Wang, Jianpeng Ma, and Wei Xu. Camel: A weakly supervised learning framework for histopathology image segmentation. In *CVPR*, pages 10682–10691, 2019. 3
- [38] Junhan Yang, Zheng Liu, Shitao Xiao, Chaozhuo Li, Defu Lian, Sanjay Agrawal, Amit Singh, Guangzhong Sun, and Xing Xie. Graphformers: Gnn-nested transformers for representation learning on textual graph. *Advances in Neural Information Processing Systems*, 34:28798–28810, 2021. 2
- [39] Jiawen Yao, Xinliang Zhu, Jitendra Jonnagaddala, Nicholas Hawkins, and Junzhou Huang. Whole slide images based cancer survival prediction using attention guided deep multiple instance learning networks. *Medical Image Analysis*, 65:101789, 2020. 1
- [40] George Zerveas, Srideepika Jayaraman, Dhaval Patel, Anuradha Bhamidipaty, and Carsten Eickhoff. A transformer-based framework for multivariate time series representation learning. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 2114–2124, 2021. 2
- [41] Hongrun Zhang, Yanda Meng, Yitian Zhao, Yihong Qiao, Xiaoyun Yang, Sarah E Coupland, and Yalin Zheng. Dtdfmil: 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*, pages 18802–18812, 2022. 1, 2, 3, 5
- [42] Xiaoxian Zhang, Sheng Huang, Yi Zhang, Xiaohong Zhang, Mingchen Gao, and Liu Chen. Dual space multiple instance representative learning for medical image classification. In *33rd British Machine Vision Conference 2022, BMVC 2022, London, UK, November 21–24, 2022*. BMVA Press, 2022. 5
- [43] Yu Zhao, Fan Yang, Yuqi Fang, Hailing Liu, et al. Predicting lymph node metastasis using histopathological images based on multiple instance learning with deep graph convolution. In *CVPR*, pages 4837–4846, 2020. 3
- [44] Yu Zhao, Zhenyu Lin, Kai Sun, Yidan Zhang, Junzhou Huang, Liansheng Wang, and Jianhua Yao. Setmil: spatial encoding transformer-based multiple instance learning for pathological image analysis. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part II*, pages 66–76. Springer, 2022. 3
- [45] Fengtao Zhou and Hao Chen. Cross-modal translation and alignment for survival analysis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 21485–21494, 2023. 5
- [46] Xinliang Zhu, Jiawen Yao, Feiyun Zhu, and Junzhou Huang. Wsisa: Making survival prediction from whole slide histopathological images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7234–7242, 2017. 2