

This ACCV 2024 paper, provided here by the Computer Vision Foundation, is the author-created version. The content of this paper is identical to the content of the officially published ACCV 2024 LNCS version of the paper as available on SpringerLink: https://link.springer.com/conference/accv

MECFormer: Multi-task Whole Slide Image Classification with Expert Consultation Network

Doanh C. Bui^{\bigcirc} and Jin Tae Kwak^{\bigcirc (\boxtimes)}

School of Electrical Engineering, Korea University, Seoul, Republic of Korea {doanhbc,jkwak}@korea.ac.kr

Abstract. Whole slide image (WSI) classification is a crucial problem for cancer diagnostics in clinics and hospitals. A WSI, acquired at gigapixel size, is commonly tiled into patches and processed by multipleinstance learning (MIL) models. Previous MIL-based models designed for this problem have only been evaluated on individual tasks for specific organs, and the ability to handle multiple tasks within a single model has not been investigated. In this study, we propose MECFormer, a generative Transformer-based model designed to handle multiple tasks within one model. To leverage the power of learning multiple tasks simultaneously and to enhance the model's effectiveness in focusing on each individual task, we introduce an Expert Consultation Network, a projection layer placed at the beginning of the Transformer-based model. Additionally, to enable flexible classification, autoregressive decoding is incorporated by a language decoder for WSI classification. Through extensive experiments on five datasets involving four different organs, one cancer classification task, and four cancer subtyping tasks, MECFormer demonstrates superior performance compared to individual state-of-theart multiple-instance learning models.

Keywords: whole slide image classification \cdot transformer \cdot multiple of expert

1 Introduction

In computational pathology, whole slide images (WSIs) have been extensively studied and investigated to provide a comprehensive understanding of tissues and diseases inside the human body and to help pathologists make accurate and reliable diagnoses in clinics [11]. Due to the gigapixel resolution of WSIs, recent studies have mainly focused on the efficient and effective processing of WSIs, which led to the development of various multiple instance learning (MIL) methods. These MIL-based methods often include two-stage procedures: 1) tiling the WSI into a bag of patches and extracting their features using an off-theshelf feature extractor; 2) learning a classifier to conduct slide-level classification. There are two primary research directions in MIL-based methods: instance-based and bag embedding-based models. Instance-based models [1, 7] provide patchlevel predictions, which are then aggregated to make the final prediction. Bag

embedding-based models [1,9,12,15,22], utilize aggregators to combine all patch features into a bag embedding, and then a bag classifier provides the final prediction. There are numerous tasks in pathology, many of which are interrelated to each other. However, most MIL-based models have been tailored and evaluated on an individual task basis under the single-task learning paradigm. This overlooks the potential to learn and leverage knowledge from multiple related tasks to resolve differing tasks, which is similar to the learning and adaptation process of human experts. One may develop a single generalizable model that can handle multiple tasks, but it is challenging for a single model to learn knowledge from several tasks simultaneously [6].



Fig. 1: Illustration of the multi-task process of MECFormer. All experts observe and process an input WSI as the corresponding expert to the target task makes the most significant contribution. A generative transformer-based model receives information from all experts and produces the predicted diagnostic term.

To address the above mentioned issues, recent efforts have developed Mixtureof-Experts (MoE) techniques [3,21], which aim to align features with the embedding space of the target task using multiple experts defined as learnable projection modules. These projections are aggregated by a weighting mechanism. MoE techniques have been applied to semantic segmentation and depth estimation tasks [3,21]; however, its utility and usefulness in WSI analysis have not been investigated yet. Moreover, most MoE methods adopt a single classification head per task that needs to be separately trained. Furthermore, a language decoder has been developed for pathology image analysis, showing its promising ability for report generation [17,24] and multi-class cancer classification [13] on patchlevel pathology images. To the best of our knowledge, no prior work has adopted the language decoder for WSI analysis.

Herein, we propose an approach for Multi-task WSI classification using an Expert Consultation Network (ECN) and a single TransFormer-based architecture, designated as MECFormer, that permits the ability to handle WSIs for various tasks. Figure 1 depicts the multi-task process of MECFormer. In

3

brief, we design MECFormer for multi-task WSI classification as a generative Transformer-based model. Given a WSI, MECFormer builds a bag of patches (or embeddings) and processes it using a Transformer-based encoder. In the encoder, we devise the first projection layer, i.e., ECN, to project the input embeddings into the model space by combining the knowledge of multiple experts via routering, task assignment, and knowledge aggregation under the guidance of the task indicator. Using the task indicator, ECN effectively gathers knowledge across multiple tasks in a controllable manner while highlighting patterns for the target task, leading to effective embeddings. Then, a Transformer-based decoder takes the knowledge aggregated embeddings along with a conditional textual input to produce the output sequence, i.e., a predicted diagnostic term, in an auto-regressive manner. We note that a single Transformer-based decoder is shared among multiple tasks. Meanwhile other MoE models adopt task-specific classification heads. Our implementation is available at ¹.

In summary, our contributions can be listed as below:

- We propose MECFormer to handle multiple well-known tasks for WSI analysis using a single unified Transformer-based model.
- We introduce ECN to efficiently and effectively learn and aggregate knowledge from multiple experts on the target task in a controllable manner.
- We adopt a generative Transformer-based decoder in MECFormer, based up cross-attention between textual and visual embeddings, to enable unified and flexible classification for multiple WSI classification tasks.
- MECFormer achieves superior performance to other state-of-the-art taskspecific MIL models on five datasets/tasks, including CAMELYON16, TCGA-BRCA, TCGA-ESCA, TCGA-NSCLC, and TCGA-RCC.

2 Related Work

2.1 Multiple Instance Learning Models for WSI analysis

Recent studies of MIL-based models commonly adopt bag embedding-based approaches. These are often built upon an attention mechanism. For example, AB-MIL [9] learns to generate a weight for each patch using the attention mechanism and then aggregates all the patch embeddings using a weighted average. CLAM [12] introduces an auxiliary task that clusters the most-attended patch embeddings as positive patches and the least-attended patch embeddings as negative patches to constrain and refine the embedding space. TransMIL [15] builds a Transformer-based architecture that approximates self-attention for computational efficiency and adopts positional encoding to retain the spatial information of the cropped patches by reshaping them into a 2-dimensional space and applying multiple learnable convolutions. DTFD-MIL [22] is designed in a two-stage manner, consisting of sub-MIL and global MIL branches, to handle over-fitting problems due to the limited number of WSIs. It randomly creates and aggregates multiple sub-bags using a sub-MIL branch, producing sub-bag embeddings. A global MIL branch makes the final prediction based on the sub-bag embeddings.

¹ https://github.com/QuIIL/MECFormer

2.2 Mixture-of-Experts Techniques for Multi-task Learning

MoE methods have recently been developed to handle multiple tasks in a single unified model. For instance, [21] proposes decomposing the feature maps into a set of multiple representative features that are processed by an expert as an MLP. A gating mechanism is devised to determine the contribution of each feature, which is used to aggregate all representative features via a weighted sum. A feature memory is adopted to capture long-range dependencies of task-specific representation. [3] proposes to use a sparse MoE layer in the Vision Transformer (ViT) [4], in which a router network is introduced for each task, along with multiple experts as MLPs. Each router network determines how many experts are employed and how they contribute to the task-specific representation. Assuming that only pre-trained models from multiple tasks are accessible, [16] utilizes task vectors [8] and a learnable router to generate weights for knowledge retrieval. These weights are combined with a base model, such as CLIP-ViT-B/32 [14], to produce features better aligned with each task.

2.3 Language Models for Pathology Analysis

Some research efforts have been made to apply and exploit language models for pathology image analysis. For instance, Zhang et al. [24] builds a network using ResNet18 to extract visual features and long short-term memory (LSTM) to generate histopathology captions. Tsuneki et al. [17] proposes to utilize EfficientNet as a visual encoder and a language decoder based on a recurrent neural network to generate reports for pathology images from gastric adenocarcinoma specimens. In an attempt to deal with multi-task pathology analysis, Nguyen et al. [13] proposes GPC for generative pathology classification, which employs a convolutional neural network as a visual encoder and a pre-trained language decoder to predict cancer categories. For slide-level classification, Bryan et al. [5] investigates the possibility of using a language decoder on WSIs. Using a dataset with image-caption pairs, the decoder is designed to learn the visual-text correspondence and to produce captions. Pairs of images and generated captions are then manually classified into a diagnostic category for evaluation.

3 Methodology

3.1 Problem Definition

In this study, we aim to develop a Transformer-based model $\mathcal{F}(\cdot)$, MECFormer, which has the ability to handle multiple tasks for WSI classification. Given a WSI \mathcal{X} and a task indicator $\mathcal{T}, \mathcal{F}(\cdot)$ produces the output that is tailored to a task of interest $\hat{\mathcal{Y}} = \mathcal{F}(\mathcal{X}, \mathcal{T})$ where $\hat{\mathcal{Y}}$ is the text output. There are T distinct tasks, each with its corresponding training dataset \mathcal{D}_t^{train} . These T datasets are merged to form $\mathcal{D}^{train} = \{\mathcal{D}_t^{train}\}_{t=1}^T$, which is used to train $\mathcal{F}(\cdot)$. Then, $\mathcal{F}(\cdot)$ is evaluated on the testing set of each t^{th} task, i.e., \mathcal{D}_t^{test} .

 $\mathbf{5}$



Fig. 2: Overview of MECFormer. MECFormer is designed as a Transformer-based generative model in an encoder-decoder manner. ECN layer is placed at the beginning of the encoder and aware of the target task along with the input bag of patch features.

3.2 MECFormer

Figure 2 illustrates the overview of MECFormer. MECFormer involves a visual encoder $\mathcal{E}(\cdot)$ and a language decoder $\mathcal{D}(\cdot)$. The role of $\mathcal{E}(\cdot)$ is to encode the bag of patch features as exploring intra-relationship between patches in an efficient and effective manner with Nyström Attention. The bag of patch features are obtained by WSI pre-processing, which is described below. $\mathcal{D}(\cdot)$ learns to align textual and visual embeddings from the encoder, and then performs auto-regressive decoding to predict a diagnostic term.

The overall procedure of MECFormer can be summarized as follows: 1) Given an input WSI \mathcal{X} and a task indicator \mathcal{T} , \mathcal{X} undergoes WSI pre-processing, producing a bag of patch features $\mathbf{x} \in \mathbb{R}^{N \times d_f}$ where N denotes the number patches and d_f refers to the feature dimension; 2) \mathbf{x} is fed into $\mathcal{E}(\cdot)$ where \mathbf{x} is embedded into the model space by an ECN projection layer $\mathcal{P}_{\text{ECN}}(\cdot)$ and processed by a series of encoding layers to produce visual token embeddings; 3) $\mathcal{D}(\cdot)$ conducts the decoding process, i.e., generating the next word, over an unknown number of K steps where each next word is conditioned on the previous word embeddings and the visual token embeddings from $\mathcal{E}(\cdot)$. The procedure can be given by:

$$\hat{y}^{(k+1)} = \mathcal{D}\Big(\mathcal{E}\big(\mathcal{P}_{\text{ECN}}(\mathbf{x},\mathcal{T})\big)\Big|\hat{y}^{(k)}\Big),\tag{1}$$

where $\hat{y}^{(k)}$ is the k^{th} prediction (or word). $\hat{y}^{(0)} = \langle BOS \rangle$, which is begin-ofsentence token to start the decoding. When $\hat{y}^{(k+1)} = \langle EOS \rangle$ is obtained, the decoding process terminates.

3.3 WSI Pre-processing

Given a WSI \mathcal{X} , a patch sampling procedure is adopted to tile \mathcal{X} into a bag of patches. In this process, we segment tissue regions on \mathcal{X} and crop patches from the segmented regions to form a bag of patches. Each patch in the bag undergoes an off-the-shelf pre-trained feature extractor $\mathcal{G}(\cdot)$, producing a bag of patch features $\mathbf{x} = \{x_i\}_{i=1}^N$ where $x_i \in \mathbb{R}^{d_f}$ and d_f is determined by $\mathcal{G}(\cdot)$. In this study, we adopt two off-the-shelf pre-trained feature extractors: CTransPath [19] and UNI [2], which are pre-trained models that are trained on large-scale pathology image datasets via self-supervised learning.

3.4 Expert Consultation Network (ECN)

To enhance the learning capability for multiple tasks, we devise an effective projection layer, $\mathcal{P}_{ECN}(\cdot)$, whose primary role is to project the input patch embeddings into the most suitable model space using the aggregated knowledge of multiple experts. Suppose that there exist multiple experts $\boldsymbol{\tau} = \{\tau_i\}_{i=1}^T$ that are designed to acquire knowledge for the corresponding tasks, and the common knowledge θ_p . During training, we learn how to obtain the knowledge of each expert and aggregate the knowledge in three stages: 1) Preliminary Consultation, 2) Expert Assignment, and 3) Expert Consultation. In Preliminary Consultation, a router $\mathcal{R}(\cdot)$ is designed to initialize T weight vectors that specify how to combine knowledge from T experts. In Expert Assignment, each expert τ_i is associated with its corresponding task by adjusting the T weight vectors. Finally, in Expert Consultation, the knowledge from T experts is aggregated using the T weight vectors and is added to the common knowledge, producing the final projection vectors to embed the bag of patches \mathbf{x} in the model space. Given an input WSI from a target t^{th} task, this process can be understood as a consultation with T experts, where the t^{th} expert has the best knowledge of the task and thus contributes the most to the final decision.

Preliminary Consultation. The router $\mathcal{R}(\cdot)$ is a stack of two linear layers. A **ReLU** activation is placed between the two layers. Given an input sequence of patch embeddings \mathbf{x} , the first layer maps the dimension size from d_f to d_{model} and the second layer produces a set of T weight vectors $\mathcal{W} = {\mathbf{w}_i}_{i=1}^T =$ $FC_2(\text{ReLU}(FC_1(\mathbf{x})))$, where FC_1 and FC_2 denote two linear layers and $\mathbf{w}_i \in \mathbb{R}^N$ is a weight vector for all N patches in a WSI.

Expert Assignment. We represent the knowledge of each expert τ_i as a learnable matrix $\mathbb{R}^{d_f \times d_{model}}$. Along with the input WSI, the task indicator \mathcal{T} specifies its task by assigning 1 to the target task and 0 to others. Using \mathcal{T} , we increase the weights of the target task \mathbf{w}_t and decrease the weights of other tasks by scaling and shifting procedures. The scaling procedure is given by:

$$\tilde{\mathcal{W}} = \{\tilde{\mathbf{w}}_i\}_{i=1}^T, \text{ where } \tilde{\mathbf{w}}_i = \frac{\exp\left(\mathbb{1}_{i=t}(\mathbf{w}_i \cdot \mathbf{1}\gamma) + \mathbb{1}_{i\neq t}(\mathbf{w}_i)\right)}{\sum_{i\neq t}^T \exp(\mathbf{w}_i) + \mathbf{w}_t \cdot \mathbf{1}\gamma}$$
(2)

where γ is a scaling hyper-parameter. In the shifting procedure, we first take the mean value of $\mathbf{\tilde{w}}_i$ to obtain a scalar weight for each task. Then, we shift the target weight to further distinguish it from other tasks. The procedure can be formulated as follows:

$$\bar{\mathbf{w}} = \{\bar{w}_i\}_{i=1}^T, \text{ where } \bar{w}_i = \frac{1}{N} \sum_{j=1}^N \tilde{w}_{ij} + \mathbb{1}_{i=t} \cdot \beta,$$
(3)

where β is a shifting hyper-parameter and \tilde{w}_{ij} is the scalar weight of the i^{th} task for the j^{th} patch in a bag of patches. In this manner, the expert τ_t is assigned the larger weight and others have smaller weights, which are determined by $\mathcal{R}(\cdot)$.

Expert Consultation. We combine the knowledge of all experts $\boldsymbol{\tau}$ using $\bar{\mathbf{w}}$, i.e., computing a weighted sum. Then, the combined knowledge is added to the common knowledge of all tasks $\theta_p \in \mathbb{R}^{d_f \times d_{model}}$, which is also a learnable matrix, producing the final projection vectors θ_p^* as follows:

$$\theta_p^* = \theta_p + \sum_{i=1}^T \tau_i \cdot \bar{w}_i. \tag{4}$$

 θ_p^* is used to project \mathbf{x} as $\mathbf{v}^{(0)} = \mathcal{P}_{\text{ECN}}(\mathbf{x}, \theta_p^*)$, aiming to produce the most suitable representation by encompassing the specialized knowledge from the target task and other tasks and the shared knowledge among all tasks. Then, $\mathbf{v}^{(0)}$ is processed by the encoder $\mathcal{E}(\cdot)$ and the decoder $\mathcal{D}(\cdot)$.

3.5 Visual Encoder

The encoder $\mathcal{E}(\cdot)$ consists of L_e encoder layers. Each layer includes a normalization layer, denoted as LN(\cdot), followed by the Nyström attention mechanism NA(\cdot) [20]. NA(\cdot) approximates self-attention operation for computational efficiency since it needs to process thousands of tokens from a WSI. A skip connection adds the input of LN(\cdot) to the output of NA(\cdot). This process for each l^{th} encoder layer can be formulated as follows:

$$\mathbf{v}^{(l)} = \mathbf{v}^{(l-1)} + \mathrm{NA}(\mathrm{LN}(\mathbf{v}^{(l-1)})), 1 \le l \le L_e,$$
(5)

where $\mathbf{v}^{(0)} = \mathcal{P}_{\text{ECN}}(\mathbf{x}, \theta_p^*).$

3.6 Language Decoder

There are L_d decoder layers in the decoder $\mathcal{D}(\cdot)$. Each decoder layer receives two inputs: one is the visual token embeddings from $\mathcal{E}(\cdot)$ and the other is the word embeddings from the previous word generated by $\mathcal{D}(\cdot)$. To produce the word embeddings, we adopt a word embedding layer $WE(\cdot, \theta_w)$ that is used to embed a set of natural words $\mathbf{s} = \{s_i\}_{i=1}^S$ (represented as token IDs) into the model space. $\theta_w \in \mathbb{R}^{N_{\text{voc}} \times d_{\text{model}}}$ and N_{voc} is the number words in the vocabulary. Then, a positional encoding PE [18] is added given by $\mathbf{h}^{(0)} = WE(\mathbf{s}) + PE$.

Each l^{th} decoder layer contains a series of masked multi-head self-attention (MHSA) and multi-head cross-attention (MHCA), of which each is preceded by $LN(\cdot)$. Masked MHSA learns the intra-relevance of the hidden tokens $\mathbf{h}^{(l)}$ without utilizing the future ground truth words during training. MHCA aligns the hidden tokens and visual token embeddings produced by $\mathcal{E}(\cdot)$. MHCA is followed by a point-wise feed-forward network (PWFF). We note that the number of tokens in $\mathbf{h}^{(l)}$ is much smaller than the number of tokens (patches) in $\mathcal{E}(\cdot)$. Hence, vanilla MHSA [18] is used for both. The process of the l^{th} decoder layer can be formulated as follows:

$$\mathbf{h}^{(l)} = \mathbf{h}^{(l-1)} + \text{Masked}_{MHSA}(\text{LN}(\mathbf{h}^{(l-1)})), \\ \mathbf{h}^{(l)} = \mathbf{h}^{(l)} + \text{MHCA}(\text{LN}(\mathbf{h}^{(l)}), \mathbf{v}), \quad , 1 \le l \le L_d.$$
(6)
$$\mathbf{h}^{(l)} = \text{PWFF}(\mathbf{h}^{(l)})$$

Following the decoder $\mathcal{D}(\cdot)$, a linear classifier $FC(\cdot, \theta_c)$ where $\theta_c \in \mathbb{R}^{d_{model} \times N_{voc}}$ is designed to produce logits $\hat{\mathbf{p}} = FC(\mathbf{h})$ where $\hat{\mathbf{p}} \in \mathbb{R}^{N_{voc}}$.

4 Experimental Results

4.1 Dataset

In this study, we consider five datasets/tasks: a breast tumor classification task (CAMELYON16) and four cancer subtyping tasks from four organs, including breast (TCGA-BRCA), esophagus (TCGA-ESCA), lung (TCGA-NSCLC), and kidney (TCGA-RCC). Figure 3 shows the distribution of the WSIs in the five datasets. To ensure the validity of our experiments, we conduct each task three times with three different trainvalidation-test splits under the same random seed. For each run, the best checkpoint on the validation set is used to evaluate the test dataset.



Fig. 3: Distribution of five datasets: CAMELYON16, TCGA-BRCA, TCGA-ESCA, TCGA-RCC and TCGA-NSCLC.

4.2 Implemental Details

The embedding size d_{model} is set to 512. The number of both encoder and decoder layers (L_e and L_d) is set to 2. For the generative decoding process, we build a vocabulary including 18 words, including <BOS> and <EOS> tokens; hence, $N_{voc} = 18$. For Nyström Attention and Multi-head Self-attention, we set the number of heads h = 8. Both scaling and shifting hyperparameters (γ and β) are set to 5. We train MECFormer for 200 epochs, using early stopping when the validation loss does not decrease for five consecutive epochs. Lookahead RAdam [23] is used as the optimizer, with the learning rate set to $1e^{-5}$ and cross-entropy loss function. All the experiments are conducted on an NVIDIA A6000 GPU.

4.3 Comparative Study

Competitors. We compare MECFormer to four other MIL-based models for WSI analysis: CLAM-MB [12], TransMIL [15], DSMIL [10], and DGCN [25]. CLAM-MB utilizes an attention mechanism and an auxiliary task on cropped patches to enhance feature representation. TransMIL proposes a learnable convolution network for positional encoding between transformer layers. DSMIL uses an instance classifier to identify the critical instance, a non-local attention mechanism to aggregate all instances along with the critical instance, and a bag classifier for the prediction. DGCN constructs a graph for a WSI based on distances between pairs of patches and adopts a graph convolution network for the classification.

Comparison setting. We compare MECFormer with four models under two settings: 1) *individual training*, where each model is trained separately for each task, and 2) *joint training*, where all five tasks are combined into a single multiclass classification. For MECFormer, we use *joint training* $+\mathcal{T}$, meaning the model is aware of the target task indicated by \mathcal{T} . Both MECFormer and the other models use two pre-trained feature extractors: CTransPath [19] and UNI [2].

Evaluation metrics. For evaluation, we use four metrics: Accuracy (Acc), F1 score (F1), Recall, and Precision. In the *joint training* setting, models may predict *invalid terms*, i.e., terms not in the ground-truth categories for a task. The overall metric is computed as $\frac{\sum_{i=1}^{N_c} m_i}{N_c + N_o}$, where m_i is F1, Precision, or Recall for each category, N_c is the number of ground-truth categories, and N_o is the number of out-of-distribution categories, further penalizing misclassifications.

4.4 Main Results

We report the classification results of MECFormer and the comparison to the state-of-the-art MIL models in Table 1 and 2.

Results on CTransPath. As shown in Table 1, using CTransPath as $\mathcal{G}(\cdot)$, MECFormer was shown to be the best performer over the four competing models under *individual training*. For CAMELYON16, TCGA-ESCA, and TCGA-NSCLC, MECFormer achieves the best performance regardless of the evaluation metrics with the improvement of $\geq 1.291\%$ Acc, ≥ 0.013 F1, ≥ 0.007 Recall, and ≥ 0.023 Precision for CAMELYON16, $\geq 0.412\%$ Acc, ≥ 0.043 F1, ≥ 0.037 Recall, and ≥ 0.036 Precision for TCGA-ESCA, and $\geq 0.412\%$ Acc and ≥ 0.004 F1,

 ≥ 0.003 Recall, and ≥ 0.003 Precision for TCGA-NSCLC. As for the remaining tasks, MECFormer was superior to other competitors except for Recall in TCGA-BRCA and Recall and Precision in TCGA-RCC. For Recall and Precision in TCGA-RCC, MECFormer was the second-best model, while it was inferior to DSMIL in Recall in TCGA-BRCA, which obtained the second-best Recall of 0.904. It is noteworthy that none of the competitors were able to outperform MECFormer in both Recall and Precision, which explains MECFormer's superior F1. In *joint training*, the four competitors obtained lower performance than their own performance in *individual training*. This is due to the invalid (or out-of-distribution) predictions made by the four competitors. These results suggest that naive multi-task learning is not specific enough to handle multiple tasks simultaneously. MECFormer, however, was not sensitive to the presence of out-of-distribution categories, resulting in a substantial performance gap; for instance, MECFormer improved F1 by 0.121~0.413, 0.121~0.329, 0.179~0.408, and 0.311~0.696 in comparison to CLAM, TransMIL, DSMIL, and DGCN, respectively, across the five tasks.

Results on UNI. Employing UNI as $\mathcal{G}(\cdot)$, we made similar observations with the above experiments with CTransPath. MECFormer demonstrated superior performance compared to the four competing models under individual training with the exception of TCGA-ESCA. In TCGA-ESCA, MECFormer achieved the second-best results for Acc, F1, and Recall, while CLAM obtained the best performance across all metrics. For CAMELYON16, TCGA-NSCLC, and TCGA-RCC, MECFormer was superior to other competitors regardless of the evaluation metrics; for instance, $\geq 0.517\%$ Acc, ≥ 0.006 F1, ≥ 0.001 Recall, and ≥ 0.010 Precision for CAMELYON16, $\geq 0.823\%$ Acc, ≥ 0.008 F1, ≥ 0.008 Recall, and ≥ 0.008 Precision for TCGA-NSCLC, and >0.383% Acc, >0.001 F1, >0.004 Recall, and >0.003 in Precision for TCGA-RCC. As for TCGA-BRCA, MECFormer improved upon other four models on Acc, F1, and Precision by >0.390%, >0.002, and ≥ 0.018 , respectively. With respect to Recall, it was inferior to CLAM, which obtained a Recall of 0.912. Under joint training, numerous invalid predictions were made by the four other models similar to the results with CTransPath. For this reason, the four competitors were substantially inferior to MECFormer; for example, F1 was dropped by $0.266 \sim 0.596$, $0.218 \sim 0.345$, $0.245 \sim 0.649$, and 0.280~0.661 for CLAM, TransMIL, DSMIL, and DGCN, respectively, in comparison to MECFormer.

4.5 Ablation Results

To deepen our understanding of MECFormer, we carried out ablation studies on two aspects: the effectiveness of 1) \mathcal{P}_{ECN} and 2) the language decoder \mathcal{D} . Here, we report the overall Acc, F1, Recall, and Precision over the five tasks. The detailed results for each task are available in the Supplementary Material.

Effectiveness of ECN. To assess the effectiveness of ECN, we analyzed three types of projection layers: 1) \mathcal{P}_1 : a single projection layer, i.e., a MLP, for all

Setting Method		Acc (%)	F1	Recall	Precision	
CAMELYON16						
Individual training	CLAM-MB [12] TransMIL [15] DSMIL [10] DGCN [25]	$\begin{array}{c} 91.731 \ (\pm 3.133) \\ 93.024 \ (\pm 1.343) \\ 90.439 \ (\pm 0.448) \\ 90.698 \ (\pm 3.379) \end{array}$	$\begin{array}{c} 0.911 \ (\pm 0.034) \\ 0.925 \ (\pm 0.014) \\ 0.897 \ (\pm 0.005) \\ 0.897 \ (\pm 0.039) \end{array}$	$\begin{array}{c c} 0.906 \ (\pm 0.033) \\ 0.919 \ (\pm 0.013) \\ 0.893 \ (\pm 0.005) \\ 0.884 \ (\pm 0.040) \end{array}$	$\begin{array}{c} 0.918 \ (\pm 0.034) \\ 0.933 \ (\pm 0.016) \\ 0.903 \ (\pm 0.005) \\ 0.921 \ (\pm 0.032) \end{array}$	
Joint training	CLAM-MB [12] TransMIL [15] DSMIL [10] DGCN [25]	$\begin{array}{c c} 92.765 \ (\pm 1.184) \\ 89.922 \ (\pm 0.775) \\ 84.496 \ (\pm 10.769) \\ 62.016 \ (\pm 0.000) \end{array}$	$ \begin{array}{c} 0.817 \ (\pm 0.168) \\ 0.793 \ (\pm 0.170) \\ 0.759 \ (\pm 0.241) \\ 0.242 \ (\pm 0.026) \end{array} $	$ \begin{vmatrix} 0.805 & (\pm 0.165) \\ 0.782 & (\pm 0.160) \\ 0.755 & (\pm 0.225) \\ 0.307 & (\pm 0.046) \end{vmatrix} $	$ \begin{array}{c c} 0.839 \ (\pm 0.176) \\ 0.819 \ (\pm 0.189) \\ 0.777 \ (\pm 0.247) \\ 0.234 \ (\pm 0.044) \end{array} $	
Joint Training + \mathcal{T}	MECFormer (ours)	$94.315 (\pm 1.614)$	$0.938 (\pm 0.018)$	$0.926 (\pm 0.020)$	$0.956~(\pm 0.014)$	
		TCGA	-BRCA			
Individual training	CLAM-MB [12] TransMIL [15] DSMIL [10] DGCN [25]	$\begin{array}{c} 94.103 \ (\pm 1.642) \\ 93.663 \ (\pm 2.674) \\ 92.813 \ (\pm 2.834) \\ 92.027 \ (\pm 1.820) \end{array}$	$\begin{array}{c} 0.896 \ (\pm 0.030) \\ 0.887 \ (\pm 0.043) \\ 0.884 \ (\pm 0.031) \\ 0.865 \ (\pm 0.026) \end{array}$	$ \begin{vmatrix} 0.896 & (\pm 0.067) \\ 0.914 & (\pm 0.070) \\ 0.904 & (\pm 0.035) \\ 0.873 & (\pm 0.054) \end{vmatrix} $	$ \begin{array}{c} 0.911 (\pm 0.053) \\ 0.871 (\pm 0.039) \\ 0.878 (\pm 0.058) \\ 0.861 (\pm 0.007) \end{array} $	
Joint training	CLAM-MB [12] TransMIL [15] DSMIL [10] DGCN [25]	$\begin{array}{c} 92.006 \ (\pm 1.527) \\ 91.957 \ (\pm 5.269) \\ 91.963 \ (\pm 3.162) \\ 88.152 \ (\pm 6.000) \end{array}$	$\begin{array}{c} 0.543 \ (\pm 0.085) \\ 0.635 \ (\pm 0.277) \\ 0.538 \ (\pm 0.076) \\ 0.512 \ (\pm 0.048) \end{array}$	$ \begin{vmatrix} 0.551 & (\pm 0.104) \\ 0.630 & (\pm 0.308) \\ 0.527 & (\pm 0.072) \\ 0.510 & (\pm 0.056) \end{vmatrix} $	$ \begin{array}{c c} 0.539 \ (\pm 0.073) \\ 0.660 \ (\pm 0.239) \\ 0.563 \ (\pm 0.088) \\ 0.524 \ (\pm 0.056) \end{array} $	
${\rm Joint \ Training} + {\cal T}$	MECFormer (ours)	94.512 (± 2.712)	$0.903 (\pm 0.047)$	$0.898 (\pm 0.077)$	$0.916~(\pm 0.026)$	
		TCGA	-ESCA			
Individual training	CLAM-MB [12] TransMIL [15] DSMIL [10] DGCN [25]	$\begin{array}{c} 89.794 \ (\pm 5.836) \\ 89.805 \ (\pm 2.193) \\ 89.805 \ (\pm 5.789) \\ 89.805 \ (\pm 5.789) \end{array}$	$\begin{array}{c} 0.895 \ (\pm 0.061) \\ 0.895 \ (\pm 0.022) \\ 0.895 \ (\pm 0.061) \\ 0.896 \ (\pm 0.060) \end{array}$	$ \begin{vmatrix} 0.903 & (\pm 0.061) \\ 0.892 & (\pm 0.020) \\ 0.894 & (\pm 0.066) \\ 0.896 & (\pm 0.064) \end{vmatrix} $	$ \begin{array}{c} 0.894 \ (\pm 0.059) \\ 0.905 \ (\pm 0.030) \\ 0.904 \ (\pm 0.050) \\ 0.899 \ (\pm 0.054) \end{array} $	
Joint training	CLAM-MB [12] TransMIL [15] DSMIL [10] DGCN [25]	$\begin{array}{c} 86.180 (\pm 9.329) \\ 91.602 (\pm 5.782) \\ 85.628 (\pm 3.575) \\ 78.994 (\pm 9.385) \end{array}$	$\begin{array}{c} 0.530 \ (\pm 0.113) \\ 0.818 \ (\pm 0.198) \\ 0.531 \ (\pm 0.102) \\ 0.487 \ (\pm 0.113) \end{array}$	$ \begin{vmatrix} 0.523 & (\pm 0.117) \\ 0.816 & (\pm 0.202) \\ 0.521 & (\pm 0.102) \\ 0.472 & (\pm 0.110) \end{vmatrix} $	$ \begin{array}{c c} 0.545 \ (\pm 0.105) \\ 0.819 \ (\pm 0.195) \\ 0.546 \ (\pm 0.094) \\ 0.532 \ (\pm 0.099) \end{array} $	
Joint Training + \mathcal{T} MECFormer (ours) 94.004 (±3.73)		94.004 (±3.739)	$0.939 (\pm 0.038)$	$0.940 (\pm 0.041)$	$0.941 (\pm 0.034)$	
TCGA-NSCLC						
Individual training	CLAM-MB [12] TransMIL [15] DSMIL [10] DGCN [25]	90.606 (± 1.843) 92.550 (± 2.874) 89.017 (± 2.876) 85.128 (± 2.055)	$\begin{array}{c} 0.906 \ (\pm 0.018) \\ 0.925 \ (\pm 0.029) \\ 0.890 \ (\pm 0.029) \\ 0.851 \ (\pm 0.020) \end{array}$	$ \begin{vmatrix} 0.907 & (\pm 0.015) \\ 0.927 & (\pm 0.027) \\ 0.893 & (\pm 0.026) \\ 0.852 & (\pm 0.019) \end{vmatrix} $	$ \begin{vmatrix} 0.911 & (\pm 0.019) \\ 0.929 & (\pm 0.027) \\ 0.895 & (\pm 0.024) \\ 0.855 & (\pm 0.021) \end{vmatrix} $	
Joint training	CLAM-MB [12] TransMIL [15] DSMIL [10] DGCN [25]	$\begin{array}{c} 87.569 \ (\pm 4.920) \\ 89.073 \ (\pm 2.060) \\ 90.194 \ (\pm 0.578) \\ 84.745 \ (\pm 1.850) \end{array}$	$\begin{array}{c} 0.516 \ (\pm 0.155) \\ 0.600 \ (\pm 0.014) \\ 0.559 \ (\pm 0.090) \\ 0.526 \ (\pm 0.075) \end{array}$	$ \begin{array}{c} 0.511 \ (\pm 0.155) \\ 0.595 \ (\pm 0.013) \\ 0.551 \ (\pm 0.088) \\ 0.518 \ (\pm 0.083) \end{array} $	$ \begin{array}{c} 0.525 \ (\pm 0.149) \\ 0.606 \ (\pm 0.014) \\ 0.569 \ (\pm 0.093) \\ 0.540 \ (\pm 0.070) \end{array} $	
${\rm Joint \ Training} + {\cal T}$	MECFormer (ours)	$92.962 (\pm 1.880)$	$0.929~(\pm 0.019)$	$0.930 (\pm 0.017)$	$0.932 \ (\pm 0.020)$	
TCGA-RCC						
Individual training	CLAM-MB [12] TransMIL [15] DSMIL [10] DGCN [25]	$\begin{array}{c} 95.788 \ (\pm 1.330) \\ 95.393 \ (\pm 3.031) \\ \textbf{96.160} \ (\pm 1.312) \\ 93.861 \ (\pm 3.303) \end{array}$	$\begin{array}{c} 0.945 \ (\pm 0.021) \\ 0.942 \ (\pm 0.033) \\ 0.950 \ (\pm 0.012) \\ 0.931 \ (\pm 0.043) \end{array}$	$ \begin{array}{c} 0.935 \ (\pm 0.011) \\ 0.936 \ (\pm 0.042) \\ 0.966 \ (\pm 0.006) \\ 0.949 \ (\pm 0.023) \end{array} $	$\begin{array}{c} 0.958 \ (\pm 0.039) \\ 0.955 \ (\pm 0.028) \\ 0.937 \ (\pm 0.022) \\ 0.921 \ (\pm 0.058) \end{array}$	
Joint training	CLAM-MB [12] TransMIL [15] DSMIL [10] DGCN [25]	$\begin{array}{c} 93.478 \ (\pm 4.023) \\ 91.936 \ (\pm 3.010) \\ 91.932 \ (\pm 1.952) \\ 92.315 \ (\pm 2.366) \end{array}$	$\begin{array}{c} 0.772 \ (\pm 0.134) \\ 0.602 \ (\pm 0.065) \\ 0.645 \ (\pm 0.082) \\ 0.646 \ (\pm 0.078) \end{array}$	$ \begin{vmatrix} 0.785 & (\pm 0.155) \\ 0.603 & (\pm 0.085) \\ 0.621 & (\pm 0.083) \\ 0.646 & (\pm 0.080) \end{vmatrix} $	$ \begin{array}{c} 0.765 \ (\pm 0.120) \\ 0.605 \ (\pm 0.049) \\ 0.621 \ (\pm 0.059) \\ 0.647 \ (\pm 0.080) \end{array} $	
Joint Training + T	MECFormer (ours)	$96.160(\pm 1.744)$	$0.957 (\pm 0.032)$	$0.960 (\pm 0.030)$	$0.957 (\pm 0.035)$	

Table 1: Classification results on five tasks using CTransPath as a feature extractor.

 The best and second-best values are highlighted in red and blue, respectively.

Setting	Method	Acc (%)	F1	Becall	Precision	
CAMELYON16						
Individual training	CLAM-MB [12]	$97.674 (\pm 1.550)$	$0.975 (\pm 0.017)$	$0.975 (\pm 0.018)$	$0.976~(\pm 0.015)$	
	TransMIL [15]	$92.506 (\pm 1.614)$	$0.917 (\pm 0.019)$	$0.901 (\pm 0.021)$	$0.946 (\pm 0.010)$	
	DSMIL [10]	$93.540 (\pm 3.663)$	$0.929 (\pm 0.041)$ 0.720 (± 0.120)	$0.918 (\pm 0.045)$ $0.724 (\pm 0.110)$	$0.948 (\pm 0.032)$	
	DGCN [20]	78.295 (±10.157)	0.759 (±0.150)	0.734 (±0.119)	0.823 (±0.113)	
	CLAM-MB [12]	$92.506 (\pm 4.408)$	$0.385 (\pm 0.089)$	$0.374 (\pm 0.092)$	$0.400 (\pm 0.081)$	
Joint training	TransMIL [15]	$95.349 (\pm 0.775)$ 72.002 (± 2.087)	$0.636 (\pm 0.005)$ 0.222 (± 0.102)	$0.627 (\pm 0.007)$ 0.227 (± 0.257)	$0.647 (\pm 0.004)$ 0.252 (±0.246)	
	DGCN [25]	$62.016(\pm 0.000)$	$0.332 (\pm 0.192)$ 0.320 (±0.109)	$0.327 (\pm 0.237)$ $0.417 (\pm 0.144)$	$0.353 (\pm 0.240)$ 0.259 (±0.088)	
		02.010 (±0.000)	0.020 (±0.105)		0.205 (±0.000)	
Joint Training $+ \gamma$	MECFormer (ours)	$98.191 (\pm 1.630)$	$0.981 (\pm 0.024)$	$0.976 (\pm 0.037)$	$0.986 (\pm 0.055)$	
		TCGA	-BRCA			
	CLAM-MB [12]	$93.700 (\pm 1.215)$	$0.895 (\pm 0.018)$	$0.912 (\pm 0.056)$	$0.886 (\pm 0.012)$	
Individual training	TransMIL [15]	$93.262 (\pm 0.913)$	$0.884 (\pm 0.019)$	$0.891 (\pm 0.064)$	$0.890 \ (\pm 0.040)$	
muividuai training	DSMIL [10]	$92.439 (\pm 1.213)$	$0.872 (\pm 0.018)$	$0.885 (\pm 0.051)$	$0.868 (\pm 0.016)$	
	DGCN [25]	$91.589 (\pm 0.820)$	$0.860 (\pm 0.016)$	$0.883 (\pm 0.076)$	$0.860 (\pm 0.044)$	
	CLAM-MB [12]	89.441 (±3.227)	$0.480 (\pm 0.081)$	$0.475 (\pm 0.061)$	$0.499 (\pm 0.120)$	
Joint training	TransMIL [15]	91.584 (±0.890)	$0.677 (\pm 0.159)$	$0.699 (\pm 0.216)$	$0.677 (\pm 0.125)$	
Joint training	DSMIL [10]	82.387 (±5.500)	$0.541 (\pm 0.077)$	$0.585 (\pm 0.074)$	$0.583 (\pm 0.064)$	
	DGCN [25]	$88.623 (\pm 3.495)$	$0.450 (\pm 0.075)$	$0.428 (\pm 0.054)$	$0.520 (\pm 0.136)$	
Joint Training + ${\cal T}$	MECFormer (ours)	$94.090 (\pm 0.448)$	$0.897 (\pm 0.005)$	$0.893 \ (\pm 0.006)$	$0.908~(\pm 0.003)$	
		TCGA	-ESCA			
	CLAM-MB [12]	$95.801 (\pm 1.087)$	$0.958(\pm 0.011)$	$0.960 (\pm 0.007)$	$0.957 (\pm 0.011)$	
* * * * * * * * *	TransMIL [15]	91.017 (±3.573)	$0.908 (\pm 0.038)$	$0.906 (\pm 0.040)$	$0.915 (\pm 0.029)$	
Individual training	DSMIL [10]	92.792 (±3.669)	$0.927 (\pm 0.036)$	$0.929 (\pm 0.031)$	$0.935 (\pm 0.032)$	
	DGCN [25]	87.446 (±3.491)	$0.871 (\pm 0.037)$	$0.870 (\pm 0.042)$	$0.878 (\pm 0.032)$	
	CLAM-MB [12]	90.985 (±4.872)	$0.510(\pm 0.053)$	$0.501 (\pm 0.056)$	$0.520 (\pm 0.051)$	
T	TransMIL [15]	$86.829(\pm 4.481)$	$0.669(\pm 0.122)$	$0.666(\pm 0.121)$	$0.688(\pm 0.145)$	
Joint training	DSMIL [10]	$86.212 (\pm 5.803)$	$0.466(\pm 0.146)$	$0.453 (\pm 0.156)$	$0.486(\pm 0.131)$	
	DGCN [25]	$80.271 (\pm 9.205)$	$0.438 (\pm 0.004)$	$0.427 (\pm 0.006)$	$0.493 (\pm 0.070)$	
Joint Training + T	MECFormer (ours)	$93.398(\pm 2.799)$	$0.934 (\pm 0.028)$	$0.939 (\pm 0.026)$	$0.934 (\pm 0.025)$	
TCGA-NSCLC						
	CLAM-MB [12]	91 457 (+3 877)	$0.914 (\pm 0.039)$	$0.916 (\pm 0.037)$	0.919(+0.035)	
	TransMIL [15]	$92.962 (\pm 2.970)$	$0.930 (\pm 0.030)$	$0.931 (\pm 0.027)$	$0.932 (\pm 0.026)$	
Individual training	DSMIL [10]	91.457 (±5.557)	$0.914 (\pm 0.055)$	$0.915 (\pm 0.054)$	$0.918 (\pm 0.056)$	
	DGCN [25]	$91.315 (\pm 2.179)$	$0.913 (\pm 0.022)$	$0.913 (\pm 0.022)$	$0.915 (\pm 0.023)$	
	CLAM-MB [12]	$91.401 (\pm 4.712)$	$0.672(\pm 0.262)$	$0.668 (\pm 0.264)$	$0.677 (\pm 0.259)$	
T • • • • •	TransMIL [15]	$90.634 (\pm 2.049)$	$0.706 (\pm 0.155)$	$0.702 (\pm 0.157)$	$0.710(\pm 0.152)$	
Joint training	DSMIL [10]	84.249 (±2.952)	$0.531 (\pm 0.304)$	$0.521 (\pm 0.313)$	$0.550 (\pm 0.304)$	
	DGCN [25]	$83.624 (\pm 3.490)$	$0.427 (\pm 0.015)$	$0.420 \ (\pm 0.017)$	$0.441 (\pm 0.009)$	
Joint Training + T	MECFormer (ours)	$93.785 (\pm 2.776)$	$0.938 (\pm 0.028)$	$0.939 (\pm 0.026)$	$0.940 (\pm 0.025)$	
TCGA-RCC						
$ $ CLAM_MR [12] $ $ 95.015 (+3.310) $ $ 0.038 (+0.045) $ $ 0.055 (+0.034) $ $ 0.027 (+0.058)						
Individual training	TransMIL [15]	$96547(\pm 2202)$	$0.958 (\pm 0.045)$ 0.958 (± 0.035)	$0.955 (\pm 0.034)$ 0.966 (± 0.030)	$0.927 (\pm 0.038)$ 0.948 (±0.045)	
	DSMIL [10]	$95.384 (\pm 1.150)$	$0.941 (\pm 0.020)$	$0.942 (\pm 0.015)$	$0.941 (\pm 0.026)$	
	DGCN [25]	$93.086 (\pm 2.272)$	$0.912 (\pm 0.024)$	$0.920 (\pm 0.014)$	$0.913 (\pm 0.049)$	
	CLAM-MB [12]	95.015 (+4.030)	0.621 (+0.109)	0.619 (+0.107)	0.624 (+0.110)	
	TransMIL [15]	$94.609(\pm 1.791)$	$0.741 (\pm 0.199)$	$0.744 (\pm 0.218)$	$0.739 (\pm 0.182)$	
Joint training	DSMIL [10]	$91.553 (\pm 6.621)$	$0.714 (\pm 0.137)$	$0.719(\pm 0.145)$	$0.717 (\pm 0.139)$	
	DGCN [25]	$91.165 (\pm 2.857)$	$0.679 (\pm 0.022)$	$0.681 (\pm 0.026)$	$0.681 (\pm 0.020)$	
Joint Training + T	MECFormer (ours)	96.930 (±1.748)	$0.959 (\pm 0.027)$	$0.970 (\pm 0.029)$	$0.951 (\pm 0.032)$	

Table 2: Classification results on five tasks using UNI as a feature extractor. The best and second-best values are highlighted in red and blue, respectively.



Fig. 4: t-SNE visualizations for four scenarios: 1) Raw features extracted by off-theshelf feature extractor $\mathcal{G}(\cdot)$, 2) \mathcal{P}_1 , 3) \mathcal{P}_T , and 4) \mathcal{P}_{ECN} . S denotes a Silhouette score.

tasks, 2) \mathcal{P}_T : a single projection layer per task, i.e., T MLPs, and 3) \mathcal{P}_{ECN} . Using each of the three projection layers, we projected the patch embeddings $\mathbf{v}^{(0)} \in \mathbb{R}^{N \times d_{model}}$, computed the mean value per channel, and obtained one representative embedding vector $v^{(0)} \in \mathbb{R}^{d_{model}}$ for each WSI. Then, t-SNE is applied to visualize the distribution of WSIs. We also computed the Silhouette score for each task. Figure 4 illustrates the visualization results and Silhouette scores. It is obvious that the embedding vectors produced by \mathcal{P}_{ECN} well separate differing tasks without any overlap between them, which is confirmed by the highest Silhouette scores of 0.762 and 0.772 for UNI and CTransPath, respectively. As for other models, there were overlaps between tasks, which can explain the generation of out-of-distribution categories.

Table 3 shows the classification results by using each of the three projection layers. Using \mathcal{P}_1 and \mathcal{P}_T , the performance was considerably dropped compared to \mathcal{P}_{ECN} ; for example, using CTransPath as \mathcal{G} , Acc, F1, Recall, and Precision decreased by $\geq 3.149\%$ Acc, ≥ 0.031 F1, ≥ 0.029 Recall, and ≥ 0.028 Precision, respectively; employing UNI as \mathcal{G} , a performance drop of $\geq 2.749\%$, ≥ 0.032 , ≥ 0.028 , and ≥ 0.033 was observed for Acc, F1, Recall, and Precision, respectively. These results validate the importance of ECN in the design of MECFormer.

Effectiveness of \mathcal{D} . To investigate the usefulness of the language decoder \mathcal{D} in a unified model for multi-task learning, we repeated the *joint learning* experiments on the five tasks without \mathcal{D} , i.e., the classifier conducts the classification for all categories across five tasks. The results are reported in Table 4. Without \mathcal{D} , the classification performance was consistently dropped regardless of the evaluation metrics and the type of \mathcal{G} . Specifically, with CTransPath as \mathcal{G} , the performance drops were 2.215 Acc, 0.025 F1, 0.022 Recall, and 0.021 Precision; with

G	Type	Acc (%)	F1	Recall	Precision
CTransPath	$\mathcal{P}_1 \ \mathcal{P}_T \ \mathcal{P}_{ extsf{eq:ecn}}$	90.438 (±0.038) 91.259 (±2.131) 94.408 (±1.879)	0.904 (±0.007) 0.905 (±0.028) 0.936 (±0.027)	0.904 (±0.014) 0.901 (±0.035) 0.933 (±0.031)	$\begin{array}{c} 0.914 \ (\pm 0.017) \\ (0.908 \pm 0.024) \\ \textbf{0.942} \ (\pm 0.022) \end{array}$
UNI	$\mathcal{P}_1 \ \mathcal{P}_T \ \mathcal{P}_{ extsf{ecn}}$	93.118 (±2.513) 92.410 (±3.092) 95.867 (±0.829)	$\begin{array}{c} 0.918 \ (\pm 0.031) \\ 0.915 \ (\pm 0.032) \\ \textbf{0.950} \ (\pm 0.011) \end{array}$	$\begin{array}{c} 0.926 \ (\pm 0.031) \\ 0.922 \ (\pm 0.040) \\ \textbf{0.954} \ (\pm 0.013) \end{array}$	$\begin{array}{c} 0.916 \ (\pm 0.031) \\ 0.913 \ (\pm 0.023) \\ \textbf{0.949} \ (\pm \textbf{0.008}) \end{array}$

Table 3: Ablation results on the three projection layer: \mathcal{P}_1 , \mathcal{P}_T , and \mathcal{P}_{ECN} .

UNI as \mathcal{G} , an Acc of 1.518, an F1 of 0.022, a Recall of 0.024, and a Precision of 0.019 decreased. These ablation results confirm the effectiveness of the language decoder \mathcal{D} in our design, which facilitates flexible and adaptable classification and improves the overall performance.

Table 4: Ablation results on the language decoder \mathcal{D} .

G	Method	Acc (%)	F1	Recall	Precision
CTransPath	$\begin{array}{c} \text{MECFormer } \texttt{w/o} \ \mathcal{D} \\ \text{MECFormer} \end{array}$	92.193 (±2.478) 94.408 (±1.879)	0.911 (±0.032) 0.936 (±0.027)	0.911 (±0.024) 0.933 (±0.031)	0.921 (±0.034) 0.942 (±0.022)
UNI	$\begin{array}{c} \text{MECFormer } \texttt{w/o} \ \mathcal{D} \\ \text{MECFormer} \end{array}$	94.169 (±2.163) 95.687 (±0.829)	0.928 (±0.030) 0.950 (±0.011)	0.930 (±0.036) 0.954 (±0.013)	0.930 (±0.025) 0.949 (±0.008)

5 Conclusion

We propose MECFormer for a unified and multi-task WSI classification. MEC-Former permits the predictions for multiple tasks across different organs using a single unified generative model, assuming that the target task is specified. To handle multiple tasks simultaneously, we introduce ECN, a projection layer that simulates expert consultation by gathering knowledge through effective gate weighting. For flexible classification, the prediction step is devised as an autoregressive decoding process using a language decoder. Through extensive experiments, we have demonstrated the superior performance of MECFormer and validated its design for multi-task WSI classification. MECFormer undoubtedly has the potential to be extended to address multi-task learning challenges in other domains.

Acknowledgments

This study was supported by a grant of the National Research Foundation of Korea (NRF) (No. 2021R1A2C2014557 and No. RS-2024-00397293) and Institute of Information & communications Technology Planning & Evaluation (IITP) (No. RS-2022-00167143).

References

- Campanella, G., Hanna, M.G., Geneslaw, L., Miraflor, A., Werneck Krauss Silva, V., Busam, K.J., Brogi, E., Reuter, V.E., Klimstra, D.S., Fuchs, T.J.: Clinicalgrade computational pathology using weakly supervised deep learning on whole slide images. Nature medicine 25(8), 1301–1309 (2019) 1, 2
- Chen, R.J., Ding, T., Lu, M.Y., Williamson, D.F., Jaume, G., Song, A.H., Chen, B., Zhang, A., Shao, D., Shaban, M., et al.: Towards a general-purpose foundation model for computational pathology. Nature Medicine **30**(3), 850–862 (2024) 6, 9
- Chen, T., Chen, X., Du, X., Rashwan, A., Yang, F., Chen, H., Wang, Z., Li, Y.: Adamv-moe: Adaptive multi-task vision mixture-of-experts. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 17346–17357 (2023) 2, 4
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020) 4
- Guevara, B.C., Marini, N., Marchesin, S., Aswolinskiy, W., Schlimbach, R.J., Podareanu, D., Ciompi, F.: Caption generation from histopathology whole-slide images using pre-trained transformers. In: Medical Imaging with Deep Learning, short paper track (2023) 4
- Guo, P., Lee, C.Y., Ulbricht, D.: Learning to branch for multi-task learning. In: International conference on machine learning. pp. 3854–3863. PMLR (2020) 2
- Hou, L., Samaras, D., Kurc, T.M., Gao, Y., Davis, J.E., Saltz, J.H.: Patch-based convolutional neural network for whole slide tissue image classification. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2424–2433 (2016) 1
- Ilharco, G., Ribeiro, M.T., Wortsman, M., Gururangan, S., Schmidt, L., Hajishirzi, H., Farhadi, A.: Editing models with task arithmetic. arXiv preprint arXiv:2212.04089 (2022) 4
- Ilse, M., Tomczak, J., Welling, M.: Attention-based deep multiple instance learning. In: International conference on machine learning. pp. 2127–2136. PMLR (2018) 2, 3
- Li, B., Li, Y., Eliceiri, K.W.: Dual-stream multiple instance learning network for whole slide image classification with self-supervised contrastive learning. In: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 14313-14323 (2021). https://doi.org/10.1109/CVPR46437.2021.01409 9, 11, 12
- Li, X., Li, C., Rahaman, M.M., Sun, H., Li, X., Wu, J., Yao, Y., Grzegorzek, M.: A comprehensive review of computer-aided whole-slide image analysis: from datasets to feature extraction, segmentation, classification and detection approaches. Artificial Intelligence Review 55(6), 4809–4878 (2022) 1
- Lu, M.Y., Williamson, D.F., Chen, T.Y., Chen, R.J., Barbieri, M., Mahmood, F.: Data-efficient and weakly supervised computational pathology on whole-slide images. Nature biomedical engineering 5(6), 555–570 (2021) 2, 3, 9, 11, 12
- Nguyen, A.T., Kwak, J.T.: Gpc: Generative and general pathology image classifier. In: Celebi, M.E., Salekin, M.S., Kim, H., Albarqouni, S., Barata, C., Halpern, A., Tschandl, P., Combalia, M., Liu, Y., Zamzmi, G., Levy, J., Rangwala, H., Reinke, A., Wynn, D., Landman, B., Jeong, W.K., Shen, Y., Deng, Z., Bakas, S., Li, X., Qin, C., Rieke, N., Roth, H., Xu, D. (eds.) Medical Image Computing and Computer

Assisted Intervention – MICCAI 2023 Workshops. pp. 203–212. Springer Nature Switzerland, Cham (2023) 2, 4

- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021) 4
- 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) 2, 3, 9, 11, 12
- Tang, A., Shen, L., Luo, Y., Yin, N., Zhang, L., Tao, D.: Merging multi-task models via weight-ensembling mixture of experts. arXiv preprint arXiv:2402.00433 (2024) 4
- Tsuneki, M., Kanavati, F.: Inference of captions from histopathological patches. In: International Conference on Medical Imaging with Deep Learning. pp. 1235–1250. PMLR (2022) 2, 4
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems **30** (2017) 7, 8
- Wang, X., Yang, S., Zhang, J., Wang, M., Zhang, J., Yang, W., Huang, J., Han, X.: Transformer-based unsupervised contrastive learning for histopathological image classification. Medical image analysis 81, 102559 (2022) 6, 9
- Xiong, Y., Zeng, Z., Chakraborty, R., Tan, M., Fung, G., Li, Y., Singh, V.: Nyströmformer: A nyström-based algorithm for approximating self-attention. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 14138–14148 (2021) 7
- Ye, H., Xu, D.: Taskexpert: Dynamically assembling multi-task representations with memorial mixture-of-experts. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 21828–21837 (2023) 2, 4
- 22. 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) 2, 3
- Zhang, M., Lucas, J., Ba, J., Hinton, G.E.: Lookahead optimizer: k steps forward, 1 step back. Advances in neural information processing systems 32 (2019) 9
- Zhang, R., Weber, C., Grossman, R., Khan, A.A.: Evaluating and interpreting caption prediction for histopathology images. In: Machine Learning for Healthcare Conference. pp. 418–435. PMLR (2020) 2, 4
- Zhao, Y., Yang, F., Fang, Y., Liu, H., Zhou, N., Zhang, J., Sun, J., Yang, S., Menze, B., Fan, X., Yao, J.: Predicting lymph node metastasis using histopathological images based on multiple instance learning with deep graph convolution. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4836–4845 (2020). https://doi.org/10.1109/CVPR42600.2020.00489 9, 11, 12