

# Towards Hierarchical Regional Transformer-based Multiple Instance Learning

Josef Cersovsky  
Bayer AG  
Berlin, Germany

josef.cersovsky@bayer.com

Dagmar Kainmueller  
Max-Delbrueck-Center for Molecular Medicine  
Berlin, Germany

dagmar.kainmueller@mdc-berlin.de

Sadegh Mohammadi  
Bayer AG  
Berlin, Germany

sadegh.mohammadi@bayer.com

Johannes Hoehne  
Berlin, Germany  
Bayer AG

johannes.hoehne@bayer.com

## Abstract

*The classification of gigapixel histopathology images with deep multiple instance learning models has become a critical task in digital pathology and precision medicine. In this work, we propose a Transformer-based multiple instance learning method that replaces the traditional learned attention mechanism with a regional, Vision Transformer inspired self-attention mechanism. We additionally propose a method that fuses regional patch information to derive slide-level predictions. We then show how this regional aggregation can be stacked to hierarchically process features on different distance levels. To increase predictive accuracy, especially for datasets with small, local morphological features, we also suggest a method to focus the image processing on high attention regions during inference. Our approach is able to significantly improve performance over the baseline on two histopathology datasets and points towards promising directions for further research.*

## 1. Introduction

Applying conventional computer vision models to gigapixel images proves to be difficult because of the large computational capacity required. The multiple instance learning (MIL) paradigm [9] allows treating an image as a non-exhaustive bag of sampled image patches, which retains the image label. A MIL-model enables generating image level predictions by either processing the patches individually or as an aggregated bag representation, thereby drastically reducing input dimensionality. A variety of MIL-based approaches working on instance-level, bag-level or a combination of both have been applied to a range of use cases [1, 4, 12]. This holds especially true in the medical domain [7, 27], where there exists an abundance of

high-resolution imaging data. MIL is commonly applied to such images by randomly sampling a bag of image patches to represent the image. However, the incomplete, sparse and often noisy nature of random patch-sampling remains a challenging hurdle and a limiting factor to predictive performance [4, 12].

Most current MIL approaches employ some form of attention-based bag aggregation [15]. They generally follow three main steps: First, the image patches are encoded into a latent space by a patch-embedder, usually a convolutional neural network (CNN). Secondly, the bag of patch-embeddings is aggregated into a single bag-level representation using a learned attention mechanism. Finally, a linear classifier derives a prediction from the bag embedding. There are three main advantages to this attention-based architecture: it greatly reduces the dimensionality by embedding patches and aggregating bags, it is invariant to the size and order of the bags and it provides high interpretability through the generated attention weights. Research has shown this approach to generalize across a variety of use cases in digital medicine [12, 14, 27].

Subsequent methods extend and improve attention-powered models applied to MIL [19]. In recent research, Transformers [28] have become the state-of-the-art for attention-based learning across many domains, including the Vision Transformer (ViT) for computer vision [10]. Input embeddings are passed through self-attention modules, where they are aggregated using a learned classification token. Transformer-based models enable a highly effective application of attention to MIL problems [24, 22, 21]. The RegionViT [5] separates patches into regions and facilitates global information exchange through a regional token. Chen et al. [6] even demonstrate a method that allows the processing of complete whole slide images by aggregating increasingly large regions in multiple steps through




Figure 1. Overview of strategies (patch/region sizes are increased for better visualization). Every depicted approach samples a total of 256 patches. (Left) Standard random global sampling, where all patches are selected randomly across the image. Captures the largest amount of variety, but does not retain any information about regionality. (Center) Regional  $2 \times 2$  sampling. This approach does pertain regional information by selecting regions, but still covers a relatively large spread of image information. (Right)  $64 \times 64$  sampling, used for multi-level aggregation. The large region-size is able to densely capture large morphological structures, but does thereby lead to a less representative sample of the full slide.

a learned classification token. However, ViT-based models can be computationally costly. They therefore usually require expensive contrastive pre-training and frozen weights for most layers. By sampling patches regionally, we reduce the required memory footprint drastically, allowing us to apply Transformers for end-to-end model training.

Various approaches also incorporate a variant of local attention based patch aggregation, reducing computational complexity and allowing the model to exploit information about patch locality [6, 29]. Konstantinov et al. [16] propose a mechanism that enriches patch embeddings by adding an aggregated representation of the patch neighborhood based on patch similarity. However, using simple dot-product attention to aggregate patches can lose valuable information during the aggregation process. We address this employing a Vision Transformer-like encoder operating on embedding-level for bag aggregation. Because using the Transformer to embed pixel-level information directly, as done in the regular ViT, leads to high computational demands, we use a CNN to derive patch embeddings.

Research has further shown that processing image patches in multiple stages can improve predictive performance [11, 25]. Myronenko et al. [21] aggregate features on different scales by downscaling the inputs in several steps using a CNN backbone. We propose a novel method of multi-level hierarchical aggregation that does not require downsampling of patches but instead enables processing of variable region sizes by iteratively aggregating increasingly large regions.

There exist several methods of pre-selecting relevant patches to guide model predictions. However, they usually require specialized architectures or additional steps during training [17, 18, 30, 23]. To leverage patch-level weights without changes to the model architecture, we instead propose a clustering based high-attention patch selection method that can be used during inference to significantly improve model performance.

## Contribution

This work aims to fuse recent developments in Transformers and MIL, by contributing a novel, regional, hierarchical Transformer architecture for efficient model training on gigapixel images. We demonstrate that our approach can leverage regional patch aggregation and the hierarchically fusion of local information on multiple levels to significantly improve over a naive global patch aggregation baseline [15]. Furthermore, we propose an easily transferable model-agnostic method for processing high-attention patches during inference, demonstrating large performance improvements on a dataset with small regions of interest. Our results showcase the potential of the proposed methodological ideas and build the groundwork for further experimentation and benchmarking.

## 2. Method

### 2.1. Multiple Instance Learning

In multiple instance learning, there exist bags of unordered instances  $\hat{X} = \{x_1, \dots, x_k\}$ . When applied to computer vision, each instance  $x_k$  is a subset of an input image  $X$ , usually corresponding to a comparatively small patch of  $X$ . Each bag is given the label  $Y$  associated with  $X$ . The labels  $y_k$  of individual instances are unknown during training. MIL models aim to determine  $Y$  by inferring the instance labels  $y_k$  or by aggregating the instances into a bag-level representation  $\mathbf{Z}$  from which  $Y$  can be deduced. In attention-based MIL, this bag representation  $\mathbf{Z}$  is based on a weighted sum of all instances inside the bag,

$$\mathbf{Z} = \sum_{k=1}^K a_k x_k, \quad (1)$$

where the attention is based on the dot-product similarity of  $x_k$  in regard to a learnable parameter  $\mathbf{C}$  with the same dimensionality as  $x_k$ :

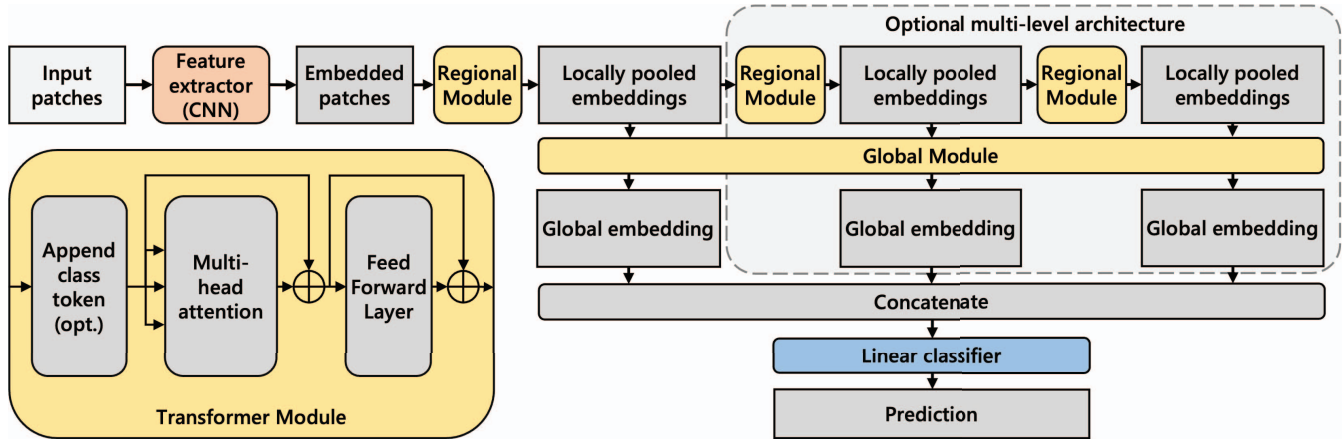


Figure 2. Proposed model architecture. A feature extraction module, usually a pretrained CNN, embeds each patch separately. Embeddings are then pooled locally by a Transformer using learnable class tokens. If multiple transformer modules are stacked, the class token is only appended in the last module. The weighted sum with regard to the class token can be considered as an embedding for the whole region. This local pooling can optionally be repeated multiple times to aggregate larger image regions. Pooled regional aggregations are finally aggregated on a slide level using a global Transformer module to generate slide-level representations. If multiple pooling-levels are used, each level’s tokens are aggregated into a separate global representation which are then concatenated. The full slide representation is fed into a linear classifier. In our experiments, weights between the regional modules are shared, while the global module uses separate weights.

$$a_k = \frac{\exp(\mathbf{C}x_k^\top)}{\sum_{i=1}^K \exp(\mathbf{C}x_i^\top)}. \quad (2)$$

We use attention-based MIL applied to patches sampled randomly from whole slide images (WSI) as done in [15] as our baseline. Each patch is a separate instance, which is represented by an embedding derived using a CNN. Note that we benchmark our method against a common, but conceptually simple baseline on two binary classification tasks. This is done on purpose, due to the paper’s focus on demonstrating novel conceptual ideas, which can be incorporated into both existing and future research. An in-depth hyperparameter tuning and evaluation of maximum performance against other state-of-the-art models and across a variety of use cases was considered out-of-scope, but remains a topic of future research.

## 2.2. Transformer based aggregation

For the Transformer model, we can formulate the attention mechanism as

$$\mathbf{Z} = \text{Attention}(Q, K, V) = \text{softmax}(QK^\top)V, \quad (3)$$

where  $Q$ ,  $K$  and  $V$  all represent the input bag  $\hat{X}$  [28]. In practice, projecting  $Q$ ,  $K$  and  $V$  linearly for  $h$  multiple times for different attention heads has been shown to improve model performance [28]. The resulting aggregates for each head are concatenated and projected back to the original input dimension  $d$ . As described in Equation 2, the bag classification is based on a learnable parameter  $C$  with the same dimensionality  $d$  as the input tokens in  $\hat{X}$ . The Vision

Transformer (ViT) [10] uses such a class token for classification. Note that the self-attention mechanism described in Equation 3 is applied to all input tokens, and therefore outputs aggregations with regard to each token in the bag. It is therefore possible to stack multiple attention layers, as the number of input tokens stays consistent. The class token  $C$  is only appended to the last layer in such cases, as it does not provide relevant information in previous layers. In our experiments, however, we only use one such attention layer and consequently only consider the  $C$ -based aggregated embedding for further processing in later stages.

The proposed model architecture is depicted in Figure 2. Instead of feeding flattened image patches into the model as done in the standard ViT approach [28], we first embed all sampled patches using a Resnet18 [13] pretrained on ImageNet [8]. The fully connected layer at the end is removed to obtain 512-dimensional image-embeddings for each patch. For our experiments, we choose a ResNet18 for its relatively small memory footprint to enable online training of the patch-embedder. The architecture and pretraining of the patch embedder can be considered tunable hyperparameters.

## 2.3. Sampling

Instead of randomly sampling patches globally, we first sample  $N$  subregions from the input image. Patches are then sampled from from each subregion separately. Each subregion has the same size as determined by the model hyperparameters. In a single level model, each subregion’s patches are aggregated into a single embedding by the regional ViT using the learned class token  $C$ . All subregional

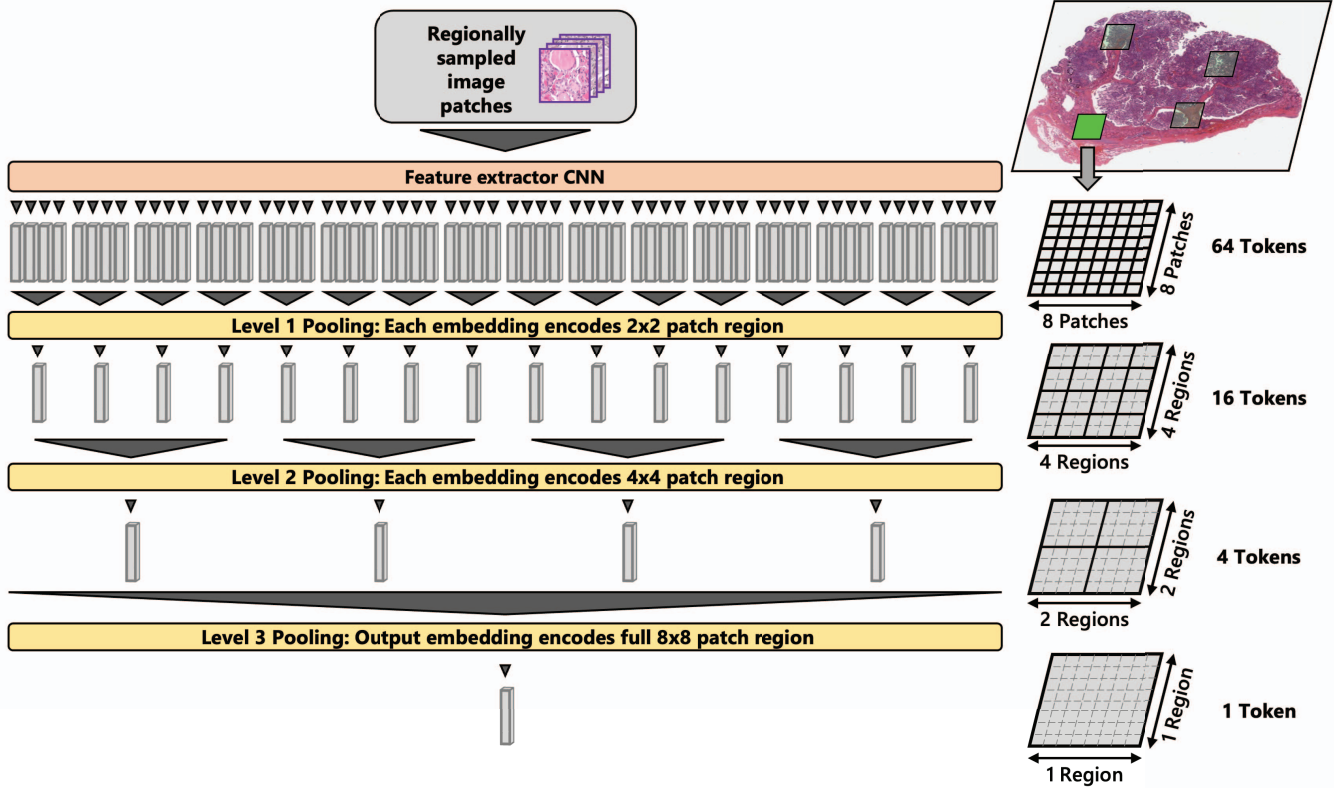


Figure 3. Illustration of embedding aggregation for a 3-level model with  $2 \times 2$  regions. Each green square represents one  $8 \times 8$  patch region. The depicted flow shows how one such region is reduced from 64 embeddings to a single aggregated representation. The size of the full region (green) depends on the number of aggregation levels (3) and the size of the aggregated regions ( $2 \times 2$ ) on each level, which are configurable hyperparameters. The pooled embeddings on every level can be passed into the global ViT module (see Figure 2).

embeddings are then passed into the global ViT to derive a slide-level representation. When using multiple levels, patches inside a subregion are pooled multiple times based on proximity. Figure 3 depicts this process for a three-level model with  $2 \times 2$  aggregation: first, all  $2 \times 2$  regions are pooled into a single embedding. This process is repeated at the next level, aggregating  $2 \times 2$  embeddings, each itself representing a  $2 \times 2$  region, into a single embedding. These new embeddings, each representing a  $4 \times 4$  patch region, are then pooled again on level 3 to derive a single embedding for the  $8 \times 8$  subregion they belong to. At every level, the regionally pooled embeddings are also fed into the global ViT to derive a slide-level representation based on each aggregation level. This aims to allow the model to consider local interactions at different distances and creates an artificial "zoom-out" that processes features at different distances, thereby enabling it to learn morphological features larger than individual patches.

We can derive the total side length ( $TSL$ ) in number of patches of a sampling region depending on the side length  $S$  of the aggregation window per level and the number of levels  $L$  as shown in Equation 4:

$$TotalSideLength = S * 2^{L-1}. \quad (4)$$

In consequence, increasing  $S$  can drastically increase the total region size: in a 3-level model ( $L = 3$ ), a base region patch width of  $S = 2$  leads to a  $TSL$  of 8 and therefore a region with 64 total patches. An increase of  $S$  by 1 already increases the  $TSL$  to 27, resulting in a region covering 729 patches. It is therefore clear that sampling on every level is required for higher values of  $L$  and especially  $S$ . In future research, adjustable aggregation window sizes per level could also be explored.

For our experiments, we use both  $2 \times 2$  ( $S = 2$ ) and  $3 \times 3$  ( $S = 3$ ) base patch regions. To enable comparability, we sample 4 patches per region in both approaches. We thereby can compare how the size of the region and the associated sparsity effects model performance. For the multi-level experiments, sampling is performed at all levels, i.e. in a 3-level  $3 \times 3 - 9 \times 9 - 27 \times 27$  model, a total of 64 patches are selected from within the  $27 \times 27$  region. We adjust the number of regions across all experiments such that there are always 256 patches sampled per whole slide image per epoch.

We use identical architectures for global and regional

Table 1. Experiment AUC results. For all patches results, all non-empty patches from the input image were passed through the model. For top patches, only selected high attention patches were used as input as described in section 2.4.

Model	Patch selection	CAMELYON16		TCGA-THCA	
		Mean	Best	Mean	Best
(1) Baseline [15]	All	0.764 ± 0.050	0.845	0.844 ± 0.010	0.855
	Top	0.865 ± 0.020	0.886	0.853 ± 0.007	0.863
<b>Single level (ours)</b>					
(3) 2 × 2	All	0.763 ± 0.035	0.833	0.852 ± 0.013	0.869
	Top	0.898 ± 0.018	0.930	0.851 ± 0.015	0.870
(4) 3 × 3	All	0.822 ± 0.039	0.873	0.855 ± 0.013	0.872
	Top	<b>0.914 ± 0.023</b>	<b>0.949</b>	0.853 ± 0.025	0.878
<b>Multi-level (ours)</b>					
(5) 2 × 2 – 4 × 4	All	0.766 ± 0.055	0.864	0.882 ± 0.012	0.898
	Top	0.867 ± 0.026	0.897	0.874 ± 0.019	0.897
(6) 2 × 2 – 4 × 4 – 8 × 8	All	0.735 ± 0.093	0.844	<b>0.891 ± 0.014</b>	<b>0.905</b>
	Top	0.823 ± 0.042	0.892	0.888 ± 0.018	0.905
(7) 3 × 3 – 9 × 9	All	0.777 ± 0.046	0.845	0.880 ± 0.010	0.893
	Top	0.832 ± 0.030	0.863	0.875 ± 0.011	0.889
(8) 3 × 3 – 9 × 9 – 27 × 27	All	0.799 ± 0.028	0.839	0.887 ± 0.017	0.905
	Top	0.857 ± 0.041	0.899	0.884 ± 0.015	0.899

ViT modules, but only share weights between the regional modules, while the global ViT retains a separate set of weights. This is done to train the model to recognize similar structures at different scales. However, the individual configurations, including number of encoders, the size of the hidden dimension, the number of attention heads and the extent to which weights are shared between modules are adjustable hyperparameters and can be tuned further. Because the model architecture is inherently designed to aggregate patches sampled from the same region and thereby implicitly introduces information about regionality, we do not use positional embeddings. Exploring the use of relative positional embeddings for large region sizes could be subject to further research.

## 2.4. Inference

During inference, the complete slide is tiled into regions and the processed by the model. This is done by first computing an embedding of all patches and then passing the embeddings through the attention mechanism, thereby significantly reducing the memory footprint.

Additionally, after deriving attention weights for all patches in the first inference pass, we apply k-means clustering with two centroids to the attention values to flexibly select high-attention patches. This step can be considered as fitting a univariate gaussian mixture model to estimate two separate distributions of high-and low attention patches respectively. We then repeat inference with all patch embeddings belonging to the low attention cluster zeroed out. This method aims to reduce the effect of noise (i.e., irrelevant patches) on the prediction, especially in cases with small regions of interest, where even a strong signal from

individual patches / regions can be diluted in the attention based aggregation process.

## 2.5. Implementation Details

Models are trained for 50 epochs using the AdamW optimizer with a learning rate of  $2e-5$  and a batch size of 2. In all experiments, 256 patches are sampled per image per epoch. All regional sampling methods sample 4 patches per lowest level. We apply random augmentations on every patch including horizontal and vertical flips, sharpness and contrast adjusting and light blurring. All experiments are repeated 5-times using different seeds for sampling, but keeping the train-validation-test split consistent.

## 2.6. Data

We evaluate our model on two publicly available datasets containing WSI of H&E-stained tissue samples. The CAMELYON16 dataset [3] depicts sentinel lymph nodes containing breast cancer metastases. We keep the train-test split predetermined by the dataset, resulting in 270 slides used for training and 129 slides used for testing. Of the train dataset, 85% is used for training and 15% for validation.

TCGA-THCA is a dataset of whole slide images of thyroid tissue provided by The Genome Cancer Atlas. The prediction of V600E mutation status of the BRAF gene mutations based on thyroid WSI has been subject of recent research [2, 14]. It contains a total of 482 tissue slides, of which 294 are labeled as BRAF positive. Of the data, 65%, 15% and 20% of cases are used for training, validation and testing respectively.




Figure 4. Comparison of two positive slides from CAMELYON16. As shown by the approximate visualizations of sampling regions, the high-attention regions of interest can be very small, extending only across few patches. Because larger aggregation regions contain a less representative sample of the full image, it is easy to see how small ROIs can be missed during sampling, hindering efficient model training. Conversely, the value of our high-attention inference method can be directly seen for cases where only few patches are relevant. This is also visible in the histograms of attention weight distributions. Both images show two peaks of high and low attention patches. However, the upper image has relatively larger tumor regions with many relevant patches and therefore receives a positive score with high confidence when predicting using all patches. For the lower image with a single small region with fewer positive patches, the prediction improves drastically when processing high attention patches.

## 2.7. Preprocessing

All images are tiled into patches with an edge size of 157 micrometers, which are then resized to a size of  $224 \times 224$  pixels. Empty patches are removed using threshold-based filtering. Luminosity standardization followed by stain normalization [26] is then applied to the patches to reduce variance between images caused by differences in staining procedures and equipment.

## 3. Results

The results of our experiments are shown in Table 1. On the CAMELYON16 dataset, we observe a notable performance increase over the baseline for both single-level regional aggregation models. Especially the  $3 \times 3$  model achieves the highest average AUC of 0.914 as well as the highest AUC observed overall of 0.949. We assume that the  $3 \times 3$  region size is well suited for the morphological size of tumors in the dataset. Additionally, the sparsity of sampling in the  $3 \times 3$  regions, compared to  $2 \times 2$ , might allow the model to better handle patches which are not directly adjacent. In contrast, multi-level approaches yield a

slightly reduced AUC when increasing the total aggregated region size, while also having larger standard deviation in results. We hypothesize the better performance of regional, Transformer based aggregation on small regions is driven mainly by the small size of regions of interest in many of the images in the dataset. As most images contain tumors spanning across only few patches, there is little potential benefit by aggregating larger regions.

On TCGA-THCA, the single level approaches do not perform notably better than the baseline. However, we see a significant improvement of results when sampling and aggregating larger regions in multiple steps. The three-level approaches in both cases outperform the respective two-level approaches. As the total region size of  $9 \times 9$  for experiment (7) is larger than that of (6), we theorize that the additional aggregation layer does help to improve predictive performance by exploiting the joint information content across a larger region. The results point towards the increased ability of the multi-level model to leverage the comparatively large tumor regions in the TCGA-THCA dataset. Furthermore, the effectiveness of 3-level approaches over 2-level approaches suggests there is potential in deriving im-

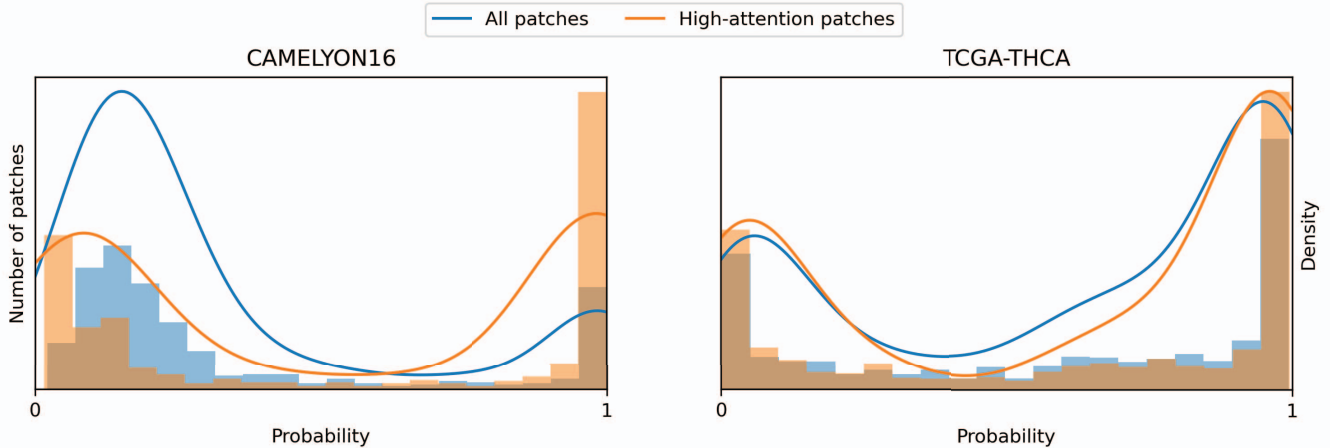


Figure 5. Predicted probabilities per slide, averaged across all seeds, using all patches and our high-attention inference strategy. Results are based on the respective best performing model (see Table 1). The high-attention strategy moves predictions closer to the extremes of 0 and 1. This is especially true for CAMELYON16. For TCGA-THCA, the effect is less pronounced, which is likely to the much larger amount of high attention patches in most images.

age representations at different “zoom” levels.

For all our experiments, we report results on the full image as well as for running inference only on high attention patches. We see significant differences in the effectiveness of this approach between the two datasets. For CAMELYON16, we find a notable improvement in AUC both for the baseline as well as for all of our approaches. As discussed above, this approach is likely particularly effective for CAMELYON16 because the signal of very small regions of interest gets diluted by large amounts of noise (i.e., patches with no relevant information about absence / presence of tumorous tissue). An example of this can be seen in Figure 4, where the prediction of the slide with a small tumor area changes drastically. Conversely, we observe no positive effect of top patch inference on TCGA-THCA, where tumor regions are much larger on average. The results highlight the potential of the proposed inference strategy, but also demonstrate that the expected gains are highly dependent on the specific morphological challenges present in any given dataset.

We were also interested whether selecting only high-attention patches would move all predictions towards positive values. As shown in Figure 5, this is not the case. The high-attention inference strategy does help to improve model certainty, both for positive and negative predictions. This indicates that the models are able to learn negative evidence which indicate the absence of tumorous tissue. More importantly, it therefore does show that performing inference on high-attention patches in negative slides does not cause the model to raise the assigned probability, underlining the robustness of our method.

## 4. Conclusion

In this work, we have demonstrated a method of regionally fusing patch embeddings to derive slide-level predictions. We show the importance of fitting the approach to the target dataset: For those with small ROIs, choosing a small, highly local aggregation method leads to best results. Datasets with large regions of interest containing complex morphological structures can be handled effectively by aggregating large regions in multi-step approaches. For both datasets, we are able to apply our method to effectively improve model performance compared to the baseline. In this work, we also suggest a simple, easily transferable method of performing inference on selected high attention patches that can significantly improve predictive results at little additional cost. Our results point towards several promising directions for further research, including an investigation into the effects of regional and global sampling with different region sizes as well as the analysis of the effects of noise on prediction results. Finally, future research could explore hyperparameter tuning and absolute performance evaluation against common state-of-the-art approaches [20, 22, 17, 6], as well as expanding these methods to use a regional and/or hierarchical aggregation approach. Similarly, we consider the application of our method to other datasets with different characteristics as well as on multi-class problems a highly interesting topic.

**Acknowledgment.** The results shown here are in whole or part based upon data generated by the TCGA Research Network: <https://www.cancer.gov/tcga>.

## References

- [1] Jaume Amores. Multiple instance classification: Review, taxonomy and comparative study. *Artificial intelligence*, 201:81–105, 2013. **1**
- [2] Deepak Anand, Kumar Yashashwi, Neeraj Kumar, Swapnil Rane, Peter H Gann, and Amit Sethi. Weakly supervised learning on unannotated h&e-stained slides predicts braf mutation in thyroid cancer with high accuracy. *The Journal of pathology*, 255(3):232–242, 2021. **5**
- [3] 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**
- [4] Marc-André Carbonneau, Veronika Cheplygina, Eric Granger, and Ghyslain Gagnon. Multiple instance learning: A survey of problem characteristics and applications. *Pattern Recognition*, 77:329–353, 2018. **1**
- [5] Chun-Fu Chen, Rameswar Panda, and Quanfu Fan. Regionvit: Regional-to-local attention for vision transformers. *arXiv preprint arXiv:2106.02689*, 2021. **1**
- [6] 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, 7**
- [7] Heather D Couture. Deep learning-based prediction of molecular tumor biomarkers from h&e: A practical review. *Journal of Personalized Medicine*, 12(12):2022, 2022. **1**
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. **3**
- [9] Thomas G Dietterich, Richard H Lathrop, and Tomás Lozano-Pérez. Solving the multiple instance problem with axis-parallel rectangles. *Artificial intelligence*, 89(1-2):31–71, 1997. **1**
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. **1, 3**
- [11] Saul Fuster, Trygve Eftestøl, and Kjersti Egan. Nested multiple instance learning with attention mechanisms. *arXiv preprint arXiv:2111.00947*, 2021. **2**
- [12] Michael Gadermayr and Maximilian Tschuchnig. Multiple instance learning for digital pathology: A review on the state-of-the-art, limitations & future potential. *arXiv preprint, arxiv.2206.04425*, 2022. **1**
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. **3**
- [14] Johannes Höhne, Jacob de Zoete, Arndt A. Schmitz, Tricia Bal, Emmanuelle di Tomaso, and Matthias Lenga. Detecting genetic alterations in braf and ntrk as oncogenic drivers in digital pathology images: towards model generalization within and across multiple thyroid cohorts. In *MIC-CAI Workshop on Computational Pathology*, pages 105–116. PMLR, 2021. **1, 5**
- [15] Maximilian Ilse, Jakub Tomczak, and Max Welling. Attention-based deep multiple instance learning. In *International conference on machine learning*, pages 2127–2136. PMLR, 2018. **1, 2, 3, 5**
- [16] Andrei V. Konstantinov and Lev V. Utkin. Multi-attention multiple instance learning. *Neural Computing and Applications*, pages 1–23, 2022. **2**
- [17] 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14318–14328, 2021. **2, 7**
- [18] Jiayun Li, Wenyuan Li, Anthony Sisk, Huihui Ye, W. Dean Wallace, William Speier, and Corey W Arnold. A multi-resolution model for histopathology image classification and localization with multiple instance learning. *Computers in biology and medicine*, 131:104253, 2021. **2**
- [19] Weijian Li, Viet-Duy Nguyen, Haofu Liao, Matt Wilder, Ke Cheng, and Jiebo Luo. Patch transformer for multi-tagging whole slide histopathology images. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 532–540. Springer, 2019. **1**
- [20] Ming Y. Lu, Drew F. K. 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. **7**
- [21] Andriy Myronenko, Ziyue Xu, Dong Yang, Holger R. Roth, and Daguang Xu. Accounting for dependencies in deep learning based multiple instance learning for whole slide imaging. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 329–338. Springer, 2021. **1, 2**
- [22] Zhuchen Shao, Hao Bian, Yang Chen, Yifeng Wang, Jian Zhang, and Xiangyang Ji. Transmil: Transformer based correlated multiple instance learning for whole slide image classification. *Advances in Neural Information Processing Systems*, 34:2136–2147, 2021. **1, 7**
- [23] Kevin Thandiackal, Boqi Chen, Pushpak Pati, Guillaume Jaume, Drew F. K. Williamson, Maria Gabrani, and Orcun Goksel. Differentiable zooming for multiple instance learning on whole-slide images. *arXiv preprint arXiv:2204.12454*, 2022. **2**
- [24] Yu Tian, Guansong Pang, Fengbei Liu, Yuyuan Liu, Chong Wang, Yuanhong Chen, Johan W. Verjans, and Gustavo Carneiro. Contrastive transformer-based multiple instance learning for weakly supervised polyp frame detection. *arXiv preprint arXiv:2203.12121*, 2022. **1**
- [25] Yanlun Tu, Houchao Lei, Wei Long, and Yang Yang. Hamil: Hierarchical aggregation-based multi-instance learn-



- ing for microscopy image classification. *arXiv preprint arXiv:2103.09764*, 2021. 2
- [26] Abhishek Vahadane, Tingying Peng, Amit Sethi, Shadi Albarqouni, Lichao Wang, Maximilian Baust, Katja Steiger, Anna Melissa Schlitter, Irene Esposito, and Nassir Navab. Structure-preserving color normalization and sparse stain separation for histological images. *IEEE transactions on medical imaging*, 35(8):1962–1971, 2016. 6
- [27] Jeroen Van der Laak, Geert Litjens, and Francesco Ciompi. Deep learning in histopathology: the path to the clinic. *Nature medicine*, 27(5):775–784, 2021. 1
- [28] 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. 1, 3
- [29] Qingfeng Wang, Ying Zhou, Jun Huang, Zhiqin Liu, Ling Li, Weiyun Xu, and Jie-Zhi Cheng. Hierarchical attention-based multiple instance learning network for patient-level lung cancer diagnosis. In *2020 IEEE international conference on bioinformatics and biomedicine (BIBM)*, pages 1156–1160. IEEE, 2020. 2
- [30] 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. 2