

# PixelCAM: Pixel Class Activation Mapping for Histology Image Classification and ROI Localization

Alexis Guichemerre<sup>1</sup>

ALEXIS.GUICHEMERRE.1@ENS.ETSMTL.CA

Soufiane Belharbi<sup>1</sup>

SOUFIANE.BELHARBI@ETSMTL.CA

Mohammadhadi Shateri<sup>1</sup>

MOHAMMADHADI.SHATERI@ETSMTL.CA

Luke McCaffrey<sup>2</sup>

LUKE.MCCAFFREY@MCGILL.CA

Eric Granger<sup>1</sup>

ERIC.GRANGER@ETSMTL.CA

<sup>1</sup> LIVIA, ILLS, Dept. of Systems Engineering, ETS Montreal, Canada

<sup>2</sup> Goodman Cancer Research, Centre, Dept. of Oncology, McGill University, Canada

**Editors:** Under Review for MIDL 2025

## Abstract

Weakly supervised object localization (WSOL) methods allow training models to classify images and localize ROIs. WSOL only requires low-cost image-class annotations, yet provides a visually interpretable classifier which is important in histology image analysis. Standard WSOL methods rely on class activation mapping (CAM) methods to produce spatial localization maps according to a single- or two-step strategy. While both strategies have made significant progress, they still face several limitations with histology images. Single-step methods can easily result in under- or over-activation due to the limited visual ROI saliency in histology images and the limited localization cues. They also face the well-known issue of asynchronous convergence between classification and localization tasks. The two-step approach is sub-optimal because it is tied to a frozen classifier, limiting the capacity for localization. Moreover, these methods also struggle when applied to out-of-distribution (OOD) datasets. In this paper, a multi-task approach for WSOL is introduced for simultaneous training of both tasks to address the asynchronous convergence problem. In particular, localization is performed in the pixel-feature space of an image encoder that is shared with classification. This allows learning discriminant features and accurate delineation of foreground/background regions to support ROI localization and image classification. We propose **PixelCAM**, a cost-effective foreground/background pixel-wise classifier in the pixel-feature space that allows for spatial object localization. Using partial-cross entropy, **PixelCAM** is trained using pixel pseudo-labels collected from a pretrained WSOL model. Both image and pixel-wise classifiers are trained simultaneously using standard gradient descent. In addition, our pixel classifier can easily be integrated into CNN- and transformer-based architectures without any modifications. Our extensive experiments<sup>1</sup> on **GlaS** and **CAMELYON16** cancer datasets show that **PixelCAM** can significantly improve classification and localization performance when integrated with different WSOL methods. Most importantly, it provides robustness on both tasks for OOD data linked to different cancer types, with large domain shifts between training and testing image data.

**Keywords:** Deep Learning, Image Classification, Visual Interpretability, Weakly Supervised Object Localization, Histology Images, Out-Of-Distribution Data.

1. <https://github.com/AlexisGuichemerreCode/PixelCAM>

## 1. Introduction

Histology image analysis remains the gold standard for diagnosing cancers of the brain (Khalsa et al., 2020), breast (Veta et al., 2014), and colon (Xu et al., 2020). However, training deep models for accurate localization of cancerous regions of interest (ROIs) requires pixel-wise annotation by pathologists which is a complex and time-consuming task, especially over Whole Slide Images (WSI). WSOL has emerged as a low-cost training approach (Zhou, 2017). Using only image class supervision, a deep model can be trained to classify an image, but also to provide spatial localization of ROIs associated with that class (Belharbi et al., 2022a; Choe et al., 2020; Murtaza et al., 2024; Rony et al., 2023). This significantly reduces the large cost of annotation and the need for dense pixel supervision. In addition, it builds interpretable classifiers (Belharbi et al., 2022b; Neto et al., 2024) which is critical in medical domains such as in histology image analysis.

Recent progress in WSOL is dominated by CAM methods (Rony et al., 2023). Several methods extract spatial activation maps in a single step (Ilse et al., 2018; Oquab et al., 2015; Wu et al., 2023), and require a module on top of a feature extractor to construct these maps. Then, a spatial pooling module extracts per-class scores used for supervised training using image-class labels. However, without any explicit guidance at the pixel level, these CAM methods can lead to poor maps due to the challenging nature of histology images where objects are often not salient (Rony et al., 2023). This can lead to under- or over-activation which creates a high false negative/positive rate at the pixel level. Moreover, training a single model may face the issue of asynchronous convergence of classification and localization tasks (Choe et al., 2020; Murtaza et al., 2024; Rony et al., 2023).

To bypass this convergence issue and the lack of pixel guidance, a recent direction has emerged in WSOL where it aims at providing explicit localization cues as pseudo-labels with dual models (Belharbi et al., 2022a,c; Murtaza et al., 2023; Wei et al., 2021; Zhang et al., 2020; Zhao et al., 2023). In particular, a per-task model is considered where the localizer is trained using pseudo-labels (Zhao et al., 2023). This leads to more parameters and training cycles. In addition, both tasks are disconnected leading to unrelated decisions in both models. A different approach (Belharbi et al., 2022a) uses a single model where a decoder is combined with a frozen classifier. While this yields good results, this architecture is limited as it is tied to frozen features at many layers using skip connections. This prevents the localizer from a better adaptation and hinders its performance.

In addition, recent work has revealed a major limitation to WSOL methods when dealing with domain shift in histology analysis (Guichemerre et al., 2024). The performance of WSOL methods is shown to decline on both tasks with out-of-distribution (OOD), limiting their real-world application. Further inspection suggests that features at the pixel level can be the root cause. Since they lack direct localization supervision, a feature encoder may provide poorly separated features concerning classes, making them vulnerable to domain shift, and class confusion.

To alleviate the aforementioned issues, we propose a novel multi-task WSOL method called PixelCAM for histology image analysis. It is based on a cost-effective foreground (FG)/background (BG) pixel-wise classifier working in the pixel-feature space, allowing for spatial object localization. It aims to explicitly learn discriminant pixel-feature representations and accurate delineation of FG and BG regions, to provide better ROI localization and



image classification. This training is achieved through localization pseudo-labels extracted from a pretrained WSOL model. In addition, such separation provides the model with robustness to OOD data. This creates a single-step multi-task framework that performs simultaneously classification and localization tasks. Multi-task training is therefore used for a better adaptation by learning rich features for both tasks, compared to tied features in a two-step approach which limits learning. This also makes both tasks cooperate to converge to a satisfying solution for the two tasks. The multi-task optimization is done using standard gradient descent by using provided image-class supervision and pixel-wise pseudo-labels extracted from a pretrained WSOL CAM model.

Our main contributions are summarized as follows. **(1)** A novel multi-task WSOL method called **PixelCAM** is proposed for histology image analysis. Multi-task optimization of both classification and localization tasks is achieved in a single step by integrating a pixel classifier at the pixel-feature level while leveraging localization pseudo-labels. **PixelCAM** alleviates the asynchronous convergence issue in WSOL, improves the performance of both tasks, and importantly, provides robustness to OOD data. Our pixel classifier is versatile in that it can easily be integrated into any CNN- or transform-based classifier architecture without modification. **(2)** We conduct extensive experiments on two public datasets for colon (**GlaS**), and breast cancer (**CAMELYON16**). Our method outperforms WSOL baseline methods on both tasks. Additionally, when dealing with large domain shifts across cancer types, our **PixelCAM** method can maintain a high level of accuracy, making it an ideal choice for OOD scenarios. We provide several ablations to further analyze our method.

## 2. Proposed Method

Let us denote by  $\mathbb{D} = \{(\mathbf{X}, y)_i\}_{i=1}^N$  a training set with  $N$  samples, where  $\mathbf{X} : \Omega \subset \mathbb{R}^2$  is an input image and  $y \in \{1, \dots, K\}$  is its global label, with  $K$  being the number of classes. Our model is composed of three parts (Fig.1): (a) feature extractor, (b) image classifier, and (c) pixel-wise classifier. (a) The feature extractor backbone  $h$  with parameter  $\theta_1$  produces a tensor spatial features  $h(\mathbf{X}) = \mathbf{F} \in \mathbb{R}^{|\Omega| \times d}$ , with depth  $d$ . For simplicity, we consider that  $\mathbf{F}$  has the same dimensions as the input image  $\mathbf{X}$  through interpolation. (b) The global image classifier head  $g$  classifies the image content, with parameters  $\theta_2$ , where  $g(\mathbf{F})_y$  is the probability of the class  $y$ . (c) The pixel-wise<sup>2</sup> classifier  $f$  which classifies the embedding  $\mathbf{F}_p \in \mathbb{R}^d$  of a pixel at location  $p$  in  $\mathbf{F}$  with parameters  $\theta_3$  into either foreground (1) or background (0) classes. This is typically a linear classifier. We refer by  $f(\mathbf{F}_p)_0$ ,  $f(\mathbf{F}_p)_1$  as the probability for the pixel at location  $p$  to be background and foreground, respectively. Classifying all locations creates the two localization maps  $\mathbf{S} = f(\mathbf{F}) \in [0, 1]^{|\Omega| \times 2}$ , where  $\mathbf{S}^0, \mathbf{S}^1$  refer to the background and foreground maps, respectively. For simplicity, we refer by  $y' \in \{0, 1\}$  as the pseudo-label of a pixel location  $p$ . Let us denote the standard cross-entropy by  $\mathbf{H}(\cdot, \cdot)$ . The total set of parameters of our model is referred to as  $\theta = \{\theta_1, \theta_2, \theta_3\}$ . When  $y, y'$  are used with  $\mathbf{H}$ , it is the one-hot encoding version that is being considered.

2. For simplicity, we refer to a location in a spatial tensor as a pixel. However, in DL models, such location typically covers more than one pixel in the image space due to the receptive field. Interpolation can be used to upscale the features to the full image size.

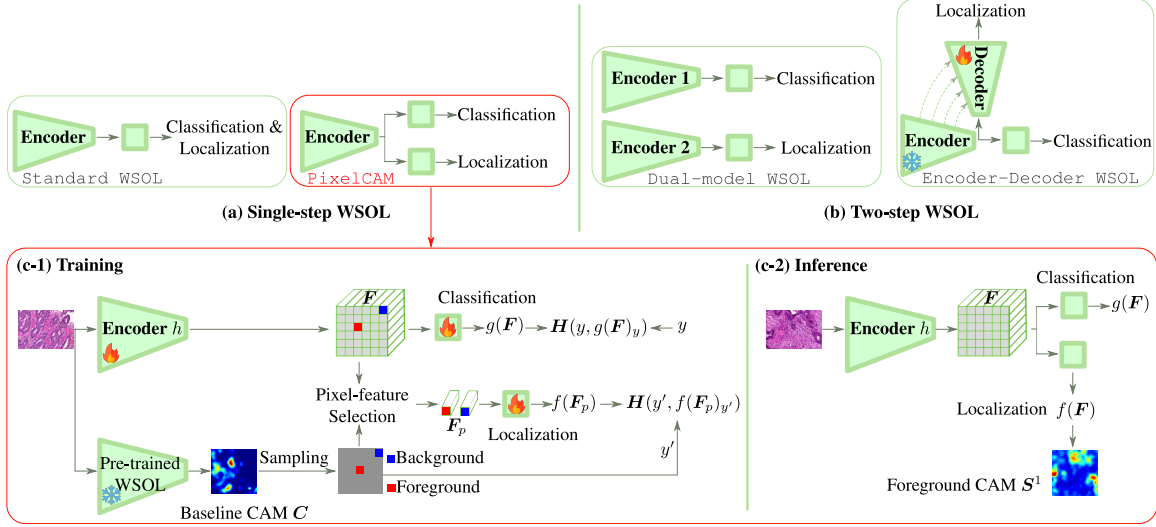


Figure 1: **Top row: WSOL learning strategies.** (a) *Single-step WSOL methods* perform classification and localization in one step. While standard WSOL uses a sequential approach using only image-class and lacks leveraging localization cues, PixelCAM relies on a multi-task training using both image-class and localization cues. (b) *Two-step WSOL methods* use a model per task or an encoder-decoder which are trained sequentially.

**Bottom row: PixelCAM training (c-1) and inference (c-2).** The training relies on pseudo-labels collected from a pretrained WSOL CAM classifier. The classification head aims to classify the extracted features  $\mathbf{F}$ . The pixel-wise classifier predicts the randomly selected locations as foreground or background. At test time, all locations are classified using the localization head  $f$  producing a localization map, while classifying the whole image using the classification head  $g$ .

To train the model for full image classification, we use standard cross-entropy over  $g$ :

$$\min_{\theta_1, \theta_2} \mathcal{H}(y, g(\mathbf{F})_y). \quad (1)$$

**Pixel-wise classifier.** Training the model using only Eq.1 builds a spatial feature  $\mathbf{F}$  that is guided only by the image classification loss. It therefore lacks awareness of localization as it does not have access to direct localization supervision cues. This can easily lead to poor pixel-wise features and confuse classes at the feature level. This is particularly true when dealing with less salient and strongly similar objects, as found in histology images. Subsequently, localization on top of these features is less reliable. Most importantly, introducing domain shift over the input image will lead to further feature confusion and failure of subsequent modules to localize and classify. To mitigate this issue, our pixel-classifier  $f$  is used to create well-separated features between foreground/background in the pixel-feature space.

Training  $f$  requires localization cues. Since the WSOL setup does not provide them, we resort to using a pretrained WSOL CAM classifier, based on CNN or transformer, to acquire

pseudo-labels. This approach is typically employed in WSOL to train a new DL model or a decoder for localization (Rony et al., 2023). However, this is cumbersome and requires another training phase with many additional parameters. To prevent this issue, our method trains a linear model to classify pixel embeddings while training the image classifier and sharing a full backbone. The number of parameters and training cycles is thereby reduced considerably.

To produce pixel-wise pseudo labels, standard WSOL methods (Rony et al., 2023) are adopted. We use the CAM  $\mathbf{C}$  of the true image class  $y$  from a pretrained WSOL CAM classifier over the same train dataset  $\mathbb{D}$ . Strong activations typically indicate foreground regions, while low activations point to background regions (Durand et al., 2017). Instead of directly fitting these regions, recent works showed that it is better to stochastically sample regions to avoid overfitting (Murtaza et al., 2023; Belharbi et al., 2022c). To sample  $n$  random foreground pixel locations, a multinomial distribution is used with  $\mathbf{C}$  as its pixel-sampling probability. Typically, this allows for sampling more frequently from high-activation locations. For background sampling,  $1 - \mathbf{C}$  is instead used as a sampling probability. Sampled pixels are collected in the sets  $\mathbb{C}^+, \mathbb{C}^-$  for foreground and background pixels, respectively. Note that sampled locations change for each image at every training step. This allows to explore different regions, and it reduces overfitting to specific regions. The sampled pseudo-labels can be directly used to train our pixel classifier, using partial cross-entropy since only part of the image space  $\Omega$  is considered at once,

$$\min_{\theta_1, \theta_3} \sum_{p \in \{\mathbb{C}^+ \cup \mathbb{C}^-\}} \mathbf{H}(y', f(\mathbf{F}_p)_{y'}) . \quad (2)$$

The pixel pseudo-label  $y'$  is collected based on the CAM of the true image class  $y$ . Therefore,  $y'$  can be perceived as an instance of the object in that image. Our pixel-classifier learns to directly recognize the class  $y$ , allowing it to perform object localization.

**Total training loss.** Unlike previous works that use a two-step approach and train different models for classification and localization, PixelCAM trains a single model with two heads simultaneously. A multi-task optimization setup is considered to train for both tasks, with most parameters being shared. To this end, we use the following composite loss,

$$\min_{\theta} \mathbf{H}(y, g(\mathbf{F})_y) + \lambda \sum_{p \in \{\mathbb{C}^+ \cup \mathbb{C}^-\}} \mathbf{H}(y', f(\mathbf{F}_p)_{y'}) , \quad (3)$$

where  $\lambda$  is a balancing coefficient. Eq.3 is minimized using standard gradient descent to perform both tasks. This mitigates the convergence issue, and allows a single training cycle. It also ensures a cooperation between localization and classification to reach a better solution for both tasks at the same time. PixelCAM does not involve a significant computation overhead. Given our lightweight pixel classifier, the number of additional parameters is negligible compared to standard DL models. Moreover, this pixel classifier is generic and can be integrated into CNN- or transformer-based models without architectural changes. Our pixel-classifier creates a boundary in the pixel feature space to well separate ROIs from noise and increase the feature’s discriminant power. This helps localization, but also the subsequent classification module as they use the same spatial features for decision. Consequently, it makes features robust to OOD data, allowing better performance of both




Figure 2: Standard WSOL setup. First column: input images from **GlaS**. Second column: Ground truth. Next columns: We display the visual CAM results for WSOL baseline without and with **PixelCAM**, respectively.

tasks in such scenario. In Section 3, we show that including this linear classifier into a standard deep WSOL model, allows us to improve its localization and classification accuracy. Moreover, it provides robustness to domain shift, a common issue in WSOL models for histology image analysis (Guichemerre et al., 2024).

### 3. Results and Discussion

#### 3.1. Experimental Methodology<sup>3</sup>

(a) *Datasets*. Our experiments are performed on two standard public datasets for WSOL in histology images. In particular, we use **GlaS** dataset (Sirinukunwattana et al., 2015) for colon cancer, and the patch version of **CAMELYON16** (Ehteshami Bejnordi et al., 2017; Rony et al., 2023) for breast cancer. In both cases, we follow the same WSOL experimental protocol used in (Belharbi et al., 2022a; Guichemerre et al., 2024; Rony et al., 2023).

(b) *Implementation details*. To assess the impact of our method, we first train a baseline WSOL method alone, and then contrast its performance when combined with our method. The pseudo-labels are extracted from the baseline model. Different recent WSOL baseline methods are used from single- and tow-step families, including DeepMIL (Ilse et al., 2018), GradCAM++ (Chattopadhyay et al., 2018), LayerCAM (Jiang et al., 2021), SAT (Wu et al., 2023), and NEGEV (Belharbi et al., 2022a). For CNN-based models, we used ResNet-50 as the backbone, and for SAT, the transformer-based model DeiT-Tiny is used. In terms of performance measures, standard WSOL metrics are used, including image classification accuracy, and pixel-localization accuracy PxAP (Choe et al., 2020; Belharbi et al., 2022c,a; Guichemerre et al., 2024). In addition, pixel-wise true/false positive/negative rates are used as well. For the localization task, we compare all WSOL methods to the full-supervised case using the U-Net model.

#### 3.2. WSOL Results

We first evaluate **PixelCAM** using the standard WSOL protocol over **GlaS** and **CAMELYON16** dataset (Rony et al., 2023). Results are reported in Tab. 1. Combining our method with a WSOL baseline often provides a significant improvement in localization accuracy. Over

3. A more detailed description of datasets and implementation are in Appendices B and C, respectively.

WSOL models	Venue	GlaS		CAMELYON16	
		PxAP	CL	PxAP	CL
DeepMIL (Ilse et al., 2018) †	<i>icml</i>	79.9	<b>100.0</b>	71.3	85.0
DeepMIL w/ NEGEV (Belharbi et al., 2022a) ★	<i>midl</i>	<b>85.9</b>	<b>100.0</b>	<u>75.6</u>	<b>89.9</b>
DeepMIL w/ PixelCAM †	-	<u>85.5</u>	<b>100.0</b>	<b>75.7</b>	<u>88.2</u>
GradCAM++ (Chattopadhyay et al., 2018) †	<i>wacv</i>	76.8	<b>100.0</b>	49.1	63.4
GradCAM++ w/ NEGEV (Belharbi et al., 2022a) ★	<i>midl</i>	<u>77.1</u>	<b>100.0</b>	<b>68.6</b>	<b>89.4</b>
GradCAM++ w/ PixelCAM †	-	<b>86.6</b>	<b>100.0</b>	<u>64.1</u>	<u>85.1</u>
LayerCAM (Jiang et al., 2021) †	<i>ieee tip</i>	<u>75.1</u>	<b>100.0</b>	33.2	84.8
LayerCAM w/ NEGEV (Belharbi et al., 2022a) ★	<i>midl</i>	73.8	<b>100.0</b>	<b>66.8</b>	<b>89.1</b>
LayerCAM w/ PixelCAM †	-	<b>83.6</b>	<b>100.0</b>	<u>66.2</u>	<b>89.1</b>
SAT (Wu et al., 2023) †	<i>iccv</i>	<u>65.9</u>	<u>98.8</u>	<u>32.8</u>	<u>83.2</u>
SAT w/ PixelCAM †	-	<b>79.1</b>	<b>100.0</b>	<b>51.2</b>	<b>87.2</b>
U-Net (Ronneberger et al., 2015)	<i>miccai</i>	95.8	n/a	81.6	n/a

Table 1: Localization (PxAP) and classification (CL) accuracy on GlaS and CAMELYON16 test sets. † refers to model with a one-step approach while ★ refers to model from two-step family.

GlaS dataset, our method improves the PxAP performance by (+5.6%, +9.8%, +8.5%) compared to the WSOL baseline methods (DeepMIL, GradCAM++, LayerCAM) respectively, while on CAMELYON16 dataset, our method gains +4.4%, +15.0%, +33.0%, respectively. Fig. 2 shows visual prediction examples. Additionally, our method improves classification performance as well. This indicates that well-separated classes at pixel-feature level in our model help improve localization but also facilitate global image classification. We further provide measures of how well are separated FG and BG pixel-features in Fig 3 on the GlaS test set. These results show that PixelCAM histogram of class separability index is shifted to the right compared to the baseline. This indicates a better class separability.




Figure 3: Histogram of foreground/background separability index for WSOL baseline methods alone vs when combined with PixelCAM on GlaS test set. The higher the value, the better the separability is. The definition of class separability is in Appendix D.

### 3.3. Domain Shift Results

Additional experiments were conducted to assess the ability of PixelCAM to train robustness models on OOD data. To this end, consider the domain shift between GlaS (colon cancer)

WSOL models	Venue	CAMELYON16 $\rightarrow$ GlaS		GlaS $\rightarrow$ CAMELYON16	
		PxAP	CL	PxAP	CL
DeepMIL (Ilse et al., 2018)	<i>icml</i>	64.5	81.2	29.0	<b>55.2</b>
DeepMIL w/ PixelCAM	-	<b>69.1</b>	<b>83.8</b>	<b>30.2</b>	52.5
GradCAM++ (Chattopadhyay et al., 2018)	<i>wacv</i>	52.9	53.7	<b>39.1</b>	52.4
GradCAM++ w/ PixelCAM	-	<b>56.2</b>	<b>71.2</b>	36.9	<b>63.3</b>
LayerCAM (Jiang et al., 2021)	<i>ieee tip</i>	59.5	77.5	30.4	51.4
LayerCAM w/ PixelCAM	-	<b>67.2</b>	<b>78.8</b>	<b>34.8</b>	<b>55.8</b>
SAT (Wu et al., 2023)	<i>iccv</i>	50.7	67.5	<b>24.3</b>	50.2
SAT w/ PixelCAM	-	<b>55.1</b>	<b>78.8</b>	22.5	<b>50.4</b>

Table 2: OOD results: Localization (PxAP) and classification (CL) accuracy on target test dataset: GlaS and CAMELYON16.

and CAMELYON16 (breast cancer) This simulates a domain adaptation problem in the source-only case, where a model is pre-trained on source domain data, and then evaluated on target data. Results are reported in Tab. 2. As shown in (Guichemerre et al., 2024), WSOL methods typically decline when facing OOD in histology data. Overall, our PixelCAM method yields better performance over both tasks and across all baselines. This is mainly due to better separated pixel-features learned by our model which make them more robust to OOD data. This is demonstrated in the appendix D by inspecting the class separability of pixel-features of the target data.

As presented in Tab. 2, using PixelCAM outperforms in general WSOL baselines alone on the case CAMELYON16  $\rightarrow$  GlaS. This improves PxAP scores by 4.6%, 3.3%, 7.7%, and 4.4% compared to DeepMIL, GradCAM++, LayerCAM, and SAT, respectively. Moreover, PixelCAM also improves image classification performance. The other case, GlaS  $\rightarrow$  CAMELYON16, is much more challenging since the train source dataset is very small, compared to the very large and difficult target set. While the overall performance is low compared to when GlaS is the target, our method still yields better localization performance in general. Moreover, classification performance is improved overall across both scenarios.

Due to space limitations, more results and interpretations are included in Appendices F and G.

## 4. Conclusion

In this paper, we propose PixelCAM, a novel multi-task single-step WSOL method. It is based on a cost-effective FG/BG pixel-wise classifier in the pixel-feature space, allowing for spatial object localization. It leverages pixel-pseudo labels cues for localization. It aims at explicitly learning discriminant pixel-features and accurately separate FG/BG regions, promoting better ROI localization and image classification. Both tasks are optimized in parallel without incurring significant computational cost nor additional parameters. Our pixel classifier is versatile. It can easily be integrated into any CNN- or transform-based classifier architecture without modification. Results on two histology datasets showed that PixelCAM can improve WSOL baseline methods over both tasks, but also make them robust to OOD data. However, there is still room to improve our method by better optimizing both tasks while reducing their mutual negative impact. In addition, dealing with OOD in WSOL scenario is still an ongoing issue, and our method still requires more improvements despite its progress.



## Acknowledgments

This research was supported in part by the Canadian Institutes of Health Research, the Natural Sciences and Engineering Research Council of Canada, and the Digital Research Alliance of Canada.

## References

- S. Belharbi, M. Pedersoli, I. Ben Ayed, L. McCaffrey, and E. Granger. Negative evidence matters in interpretable histology image classification. In *MIDL*, 2022a.
- S. Belharbi, J. Rony, J. Dolz, I. Ben Ayed, L. McCaffrey, and E. Granger. Deep interpretable classification and weakly-supervised segmentation of histology images via max-min uncertainty. *TMI*, 41:702–714, 2022b.
- S. Belharbi, A. Sarraf, M. Pedersoli, I. Ben Ayed, L. McCaffrey, and E. Granger. F-CAM: Full resolution class activation maps via guided parametric upscaling. In *WACV*, 2022c.
- S. Belharbi, I. Ben Ayed, L. McCaffrey, and E. Granger. Tcam: Temporal class activation maps for object localization in weakly-labeled unconstrained videos. In *WACV*, 2023.
- S. Belharbi, S. Murtaza, M. Pedersoli, I. Ben Ayed, L. McCaffrey, and E. Granger. Colocam: Class activation mapping for object co-localization in weakly-labeled unconstrained videos. *PR*, 162:111358, 2025.
- A. Chattopadhyay, A. Sarkar, P. Howlader, and V. N. Balasubramanian. Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. In *WACV*, pages 839–847, 2018.
- J. Choe, S. J. Oh, S. Lee, S. Chun, Z. Akata, and H. Shim. Evaluating weakly supervised object localization methods right. In *CVPR*, 2020.
- R.O. Duda, P.E. Hart, and D.G. Stork. *Pattern Classification*. Wiley-Interscience, 2 edition, November 2000. ISBN 0471056693.
- T. Durand, T. Mordan, N. Thome, and M. Cord. Wildcat: Weakly supervised learning of deep convnets for image classification, pointwise localization and segmentation. In *CVPR*, 2017.
- B. Ehteshami Bejnordi, M. Veta, P. Johannes van Diest, 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.
- W. Gao, F. Wan, X. Pan, Z. Peng, Q. Tian, Z. Han, B. Zhou, and Q. Ye. TS-CAM: token semantic coupled attention map for weakly supervised object localization. In *ICCV*, 2021.
- A. Guichemerre, S. Belharbi, T. Mayet, S. Murtaza, P. Shamsolmoali, L. McCaffrey, and E. Granger. Source-free domain adaptation of weakly-supervised object localization models for histology. In *CVPRw*, 2024.

- M. Ilse, J. M. Tomczak, and M. Welling. Attention-based deep multiple instance learning. In *ICML*, 2018.
- P.-T. Jiang, C.-B. Zhang, Q. Hou, M.-M. Cheng, and Y. Wei. Layercam: Exploring hierarchical class activation maps for localization. *TMI*, 30:5875–5888, 2021.
- S.S.S. Khalsa, T.C. Hollon, A. Adapa, E. Urias, S. Srinivasan, N. Jairath, J. Szczepanski, P. Ouillette, S. Camelo-Piragua, and D.A. Orringer. Automated histologic diagnosis of cns tumors with machine learning. *CNS oncology*, 9(2):CNS56, 2020.
- A. Krizhevsky, I. Sutskever, and G.E. Hinton. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, 2012.
- S. Murtaza, S. Belharbi, M. Pedersoli, A. Sarraf, and E. Granger. Dips: Discriminative pseudo-label sampling with self-supervised transformers for weakly supervised object localization. *IMAVIS*, page 104838, 2023.
- S. Murtaza, S. Belharbi, M. Pedersoli, and E. Granger. A realistic protocol for evaluation of weakly supervised object localization. In *WACV*, 2024.
- P. C. Neto, D. Montezuma, S. P. Oliveira, D. Oliveira, J. Fraga, A. Monteiro, J. Monteiro, L. Ribeiro, S. Gonçalves, S. Reinhard, I. Zlobec, I. M. Pinto, and J. S. Cardoso. An interpretable machine learning system for colorectal cancer diagnosis from pathology slides. *NPJ precision oncology*, 8(1):56, 2024.
- M. Oquab, L. Bottou, I. Laptev, and J. Sivic. Is object localization for free?-weakly-supervised learning with convolutional neural networks. In *CVPR*, 2015.
- O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015.
- J. Rony, S. Belharbi, J. Dolz, I. Ben Ayed, L. McCaffrey, and E. Granger. Deep weakly-supervised learning methods for classification and localization in histology images: A survey. *MLBI*, 2:96–150, 2023.
- R.R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. *IJCV*, 128(2):336–359, 2020.
- K. Sirinukunwattana, D.RJ Snead, and N.M. Rajpoot. A stochastic polygons model for glandular structures in colon histology images. *TMI*, 34(11):2366–2378, 2015.
- M. Veta, J.P.W. Pluim, P. J. Van Diest, and M.A. Viergever. Breast cancer histopathology image analysis: A review. *IEEE Transactions on Biomedical Engineering*, 61:1400–1411, 2014.
- J. Wei, Q. Wang, Z. Li, S. Wang, S. K. Zhou, and S. Cui. Shallow feature matters for weakly supervised object localization. In *CVPR*, 2021.
- P. Wu, W. Zhai, Y. Cao, J. Luo, and Z.-J. Zha. Spatial-aware token for weakly supervised object localization. In *ICCV*, 2023.

- L. Xu, B. Walker, P.-I. Liang, Y. Tong, C. Xu, Y.C. Su, and A. Karsan. Colorectal cancer detection based on deep learning. *Journal of Pathology Informatics*, 11, 2020.
- C.-L. Zhang, Y.-H. Cao, and J. Wu. Rethinking the route towards weakly supervised object localization. In *CVPR*, 2020.
- Y. Zhao, Q. Ye, W. Wu, C. Shen, and F. Wan. Generative prompt model for weakly supervised object localization. In *CVPR*, 2023.
- B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. Learning deep features for discriminative localization. In *CVPR*, 2016.
- Z.-H. Zhou. A brief introduction to weakly supervised learning. *National Science Review*, 5(1):44–53, 2017.
- L. Zhu, Q. She, Q. Chen, Q. Ren, and Y. Lu. Boosting weakly supervised object localization and segmentation with domain adaption. *IEEE TPAMI*, 2024.

This supplementary material contains the following points:

1. Related work
2. Datasets details
3. Training details for standard WSOL and OOD scenarios
4. Class separability index
5. More localization performance analysis
6. Ablations:
  - (a) Impact of PixelCAM depth
  - (b) Impact of model selection of CAM pre-trained model
  - (c) Impact of pixel sampling technique
  - (d) Impact of the number of sampled pixels  $n$
  - (e) Impact of  $\lambda$
7. More visualizations results:
  - (a) Standard WSOL setup
  - (b) Over Target set with OOD setup

## Appendix A. Related Work

WSOL has emerged as a low-cost training setup to localize objects while classifying an image content (Choe et al., 2020; Murtaza et al., 2024; Zhou, 2017). WSOL has also been extended to videos (Belharbi et al., 2023, 2025). Moreover, there has been a recent focus on designing WSOL methods for histology image analysis (Rony et al., 2023). This allows us to reduce the large annotation cost, in addition to building visually interpretable classifiers. **Single-step WSOL.** Several approaches address the WSOL problem in a single step (Choe et al., 2020; Murtaza et al., 2024; Rony et al., 2023) where a single model is trained to do both tasks, classification and localization. Usually, this is achieved by using a localization head followed by a spatial pooling layer to extract per-class probability. Early (Ilse et al., 2018; Selvaraju et al., 2020; Zhou et al., 2016) but also recent WSOL works (Wu et al., 2023; Zhu et al., 2024; Gao et al., 2021) follow this strategy. Although this has brought significant improvements into the field, these models still face several limitations. Since only image-class labels are used as supervision without any localization cues, CAM localization can have poor performance, more so when dealing with less salient objects such as in histology images. Recent work (Rony et al., 2023) showed that without localization cues over histology images, these models often lead to highly unbalanced localization. This manifests either by under- or over-activations leading to high false negative/positive rates. In addition, this approach faces a great challenge in model selection as both tasks proceed in an asynchronous convergence (Choe et al., 2020; Murtaza et al., 2024; Rony et al., 2023). For instance, the model selected for best localization often yields poor classification.

**Two-step WSOL.** This strategy has emerged as a new line of research. In particular, a dual-model approach is used in combination with the usage of localization cues under the form of pseudo-labels (Murtaza et al., 2024; Rony et al., 2023). This provides direct guidance for the localization task, bypassing the issue of task convergence. Several works simply create a dedicated model per task: one for classification and another for localization (Wei et al., 2021; Zhang et al., 2020; Zhao et al., 2023). This leads to better performance since tasks are divided between two large models. However, this is a cumbersome approach as it drastically increases the number of parameters and training cycles. Most importantly, the localization CAMs are completely disconnected from the classification decision, making this strategy unreliable and less interpretable. A parallel direction overcomes this issue by using a decoder to act as a localizer on top of the classifier (Belharbi et al., 2022a,c; Murtaza et al., 2023), in a similar way to a U-Net architecture (Ronneberger et al., 2015). This creates a direct relation between both tasks. However, the decoder has a limited learning capacity since it is tied to a frozen encoder at several layers via skip connections. This prevents the localizer from a better adaptation, and it limits its performance leading to a sub-optimal solution. In addition, it adds a significant number of parameters to the classifier.

**WSOL vs OOD.** In addition to these limitations, WSOL methods face challenges when dealing with domain shift which degrades the performance of both tasks (Guichemerre et al., 2024). Such shift is common in histology due to variations in stains, objects’ structure, microscope type, and imaging centers. Further analysis shows that this issue could be rooted in the pixel features of the image encoder. Both classification and localization tasks depend heavily on these spatial features. Less flexible and poorly discriminant pixel features can lead to poor performance in both tasks since they both rely on these embeddings.

In summary, while existing WSOL methods have achieved great progress they still face several limitations. Single-step WSOL lacks the leverage of localization cues making them vulnerable to wrong localization when dealing with complex and less salient data such as histology images. This adds to the well known issue of asynchronous convergence of localization and classification tasks. On the other hand, two-step methods are cumbersome in terms of training cycles, and number of parameters. In addition, either both tasks are disconnected leading to misaligned decisions or the localization has a tied capacity due to frozen backbone. Our method comes as simple yet efficient alternative. It performs WSOL in a single-step within a multi-task framework where both tasks are optimized simultaneously. This leads to a single training cycle and facilitate convergence issue. It allows sharing parameters between both tasks making it parameter-efficient. But also, it can leverage localization cues such as pseudo-labels. Finally, well separating features at pixel level makes our model robust to OOD data.

## Appendix B. Datasets

**GlaS.** The GlaS dataset is used for the diagnosis of colon cancer. The dataset consists of 165 images from 16 Hematoxylin and Eosin (H&E) and includes labels at both pixel-level and image-level (benign or malign). The dataset consists of 67 images for training, 18 for validation, and 80 for testing. We use the same protocol as in (Rony et al., 2023; Guichemerre et al., 2024; Belharbi et al., 2022a). Similarly, we use 3 samples per class with full supervision in validation set for model selection for localization (B-LOC).

**CAMELYON16.** A patch-based benchmark (Rony et al., 2023) is extracted from the CAMELYON16 dataset that contains 399 Whole slide images categorized into two classes (normal and metastatic) for the detection of breast cancer metastases in H&E-stained tissue sections of sentinel lymph nodes. Patch extraction of size  $512 \times 512$  follows a protocol established by (Rony et al., 2023; Guichemerre et al., 2024; Belharbi et al., 2022a) to obtain patches with annotations at the image and pixel level. The dataset contains a total of 48870 images, including 24348 for training, 8850 for validation, and 15664 for testing. From the validation dataset, 6 examples per class are randomly selected to be fully supervised to perform model selection for localization (B-LOC) similarly to (Rony et al., 2023; Guichemerre et al., 2024; Belharbi et al., 2022a).

## Appendix C. WSOL Training Details

For pretraining a WSOL baseline method on the data, we use the same setup as in (Guichemerre et al., 2024). The first part of the model training is defined using the backbone trained on ImageNet (Krizhevsky et al., 2012). The training is done using SGD with a batch size of 32 (Guichemerre et al., 2024). For the GLaS dataset, training is performed over 1000 epochs, and 20 epochs for CAMELYON16. A weight decay of  $10^{-4}$  is also used. During training, images are resized to  $256 \times 256$ , then randomly cropped to  $224 \times 224$ . A hyperparameter search was conducted for the learning rate parameters among the values  $\{0.0001, 0.001, 0.01\}$ , and its decaying factor among  $\{0.1, 0.4, 0.9\}$  following (Guichemerre et al., 2024). In the second phase of training PixelCAM, we use the CAMs generated by the previous method and continue the training using the same setup as defined previously. For the OOD scenario, the source model is trained using the same setup on a source dataset and evaluated on a target dataset.

## Appendix D. Class Separability Index

In this section, we present the class separability index between foreground (FG) and background (BG) classes at pixel-feature level over test set of both datasets GLaS and CAMELYON16. To measure how well FG/BG classes are separated in the feature space of pixels, we resort to the class separability index (Duda et al., 2000). The class separability,  $J$ , is based on the Within-class scatter matrix ( $S_W$ ) and the Between-class scatter matrix ( $S_B$ ) of pixel-features, defined as follows,

$$\mathbf{S}_W = \sum_{i=1}^c \left[ \sum_{j=1}^{n_i} (\mathbf{x}_{i,j} - \mathbf{m}_i)(\mathbf{x}_{i,j} - \mathbf{m}_i)^\top \right], \quad (4)$$

$$\mathbf{S}_B = \sum_{i=1}^c n_i (\mathbf{m}_i - \mathbf{m})(\mathbf{m}_i - \mathbf{m})^\top, \quad (5)$$

where  $c$  is the number of class (in our case  $c = 2$  for foreground and background). Note that this index  $J$  is computed over a single image using the pixel-labels. Let  $n_i (i = 0, \dots, c)$  the number of pixels in the  $i$ -th class. The feature vector  $\mathbf{x}_{i,j} \in \mathbb{R}^d$  denotes the  $j$ -th pixel of the  $i$ -th class. It represents the vector  $\mathbf{F}_p$  for the pixel  $p$  in the main paper. The mean



vector of the  $i$ -th class is given by  $\mathbf{m}_i$ , while  $\mathbf{x}$  represents the mean vector computed over all feature vectors. The class separability denoted  $\mathbf{J}$  is defined as follows (Duda et al., 2000),

$$\mathbf{J} = \frac{\mathbf{S}_B}{\mathbf{S}_W}. \quad (6)$$

In Table 3 and 4, we provide the average class separability over normal and cancer classes and the overview over the entire dataset. As we can observe, PixelCAM improves the class separability between FG and BG of WSOL baseline methods on both dataset except for GradCAM++ on CAMELYON16. This separability is explained by the high number of images with a low separability as illustrated in Fig 4. The improvement of PixelCAM in the OOD scenario compared to the WSOL baseline methods such as DeepMIL and LayerCAM on GlaS and CAMELYON16 can be attributed to better separability as shown in Figs 5 and 6. We provide examples of features separability at the pixel level in section G.

WSOL models	GlaS			CAMELYON16		
	Normal	Cancer	All	Normal	Cancer	All
DeepMil (Ise et al., 2018)(icml)	0.13	0.07	0.10	-	0.11	0.11
DeepMil w/ PixelCAM	<b>0.21</b>	<b>0.09</b>	<b>0.14</b>	-	<b>0.17</b>	<b>0.17</b>
GradCAM++ (Chattopadhyay et al., 2018)(wacv)	0.22	0.13	0.17	-	<b>0.22</b>	<b>0.22</b>
GradCAM++ w/ PixelCAM	<b>0.25</b>	<b>0.15</b>	<b>0.20</b>	-	0.16	0.16
LayerCAM (Jiang et al., 2021)(ieee tip)	<b>0.17</b>	0.04	0.10	-	0.07	0.07
LayerCAM w/ PixelCAM	0.14	<b>0.12</b>	<b>0.13</b>	-	<b>0.17</b>	<b>0.17</b>
SAT (Wu et al., 2023)(iccv)	0.21	0.06	0.13	-	0.17	0.17
SAT w/ PixelCAM	<b>0.25</b>	<b>0.10</b>	<b>0.17</b>	-	<b>0.22</b>	<b>0.22</b>

Table 3: Standard WSOL setup: Average class separability index  $J$  between FG/BG pixel-features on test set GlaS and CAMELYON16 for WSOL baselines with and without PixelCAM. The higher  $J$  is, the more classes are separated.

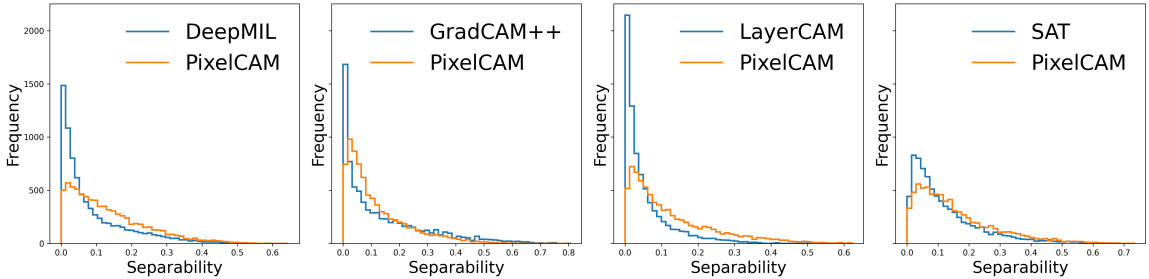


Figure 4: Standard WSOL setup: Histogram of class separability index  $J$  between FG/BG pixel-features on test set CAMELYON16 for WSOL baselines with and without PixelCAM. The higher  $J$  is, the more classes are separated.

WSOL models	CAMELYON16 $\rightarrow$ GlaS			GlaS $\rightarrow$ CAMELYON16		
	Normal	Cancer	All	Normal	Cancer	All
DeepMil (Ise et al., 2018) ( <i>icml</i> )	0.05	0.11	0.08	-	<b>0.05</b>	<b>0.05</b>
DeepMil w/ PixelCAM	<b>0.06</b>	<b>0.13</b>	<b>0.10</b>	-	<b>0.05</b>	<b>0.05</b>
GradCAM++ (Chattopadhyay et al., 2018) ( <i>wacv</i> )	<b>0.08</b>	<b>0.19</b>	<b>0.14</b>	-	<b>0.08</b>	<b>0.08</b>
GradCAM++ w/ PixelCAM	0.05	0.09	0.07	-	0.07	0.07
LayerCAM (Jiang et al., 2021) ( <i>ieee tip</i> )	0.02	<b>0.13</b>	0.08	-	0.04	0.04
LayerCAM w/ PixelCAM	<b>0.06</b>	0.11	<b>0.09</b>	-	<b>0.09</b>	<b>0.09</b>
SAT (Wu et al., 2023) ( <i>iccv</i> )	0.11	0.09	0.10	-	0.08	0.08
SAT w/ PixelCAM	<b>0.12</b>	<b>0.12</b>	<b>0.12</b>	-	<b>0.09</b>	<b>0.09</b>

Table 4: OOD setup: Average class separability index  $J$  between FG/BG pixel-features on target test set GlaS and CAMELYON16 for both OOD cases: CAMELYON16  $\rightarrow$  GlaS and GlaS  $\rightarrow$  CAMELYON16 for WSOL baselines with and without PixelCAM. The higher  $J$  is, the more classes are separated.




Figure 5: OOD setup: Histogram of class separability index  $J$  between FG/BG pixel-features on target test set GlaS for the OOD case: CAMELYON16  $\rightarrow$  GlaS for WSOL baselines with and without PixelCAM. The higher  $J$  is, the more classes are separated.




Figure 6: OOD setup: Histogram of class separability index  $J$  between FG/BG pixel-features on target test set CAMELYON16 for the OOD case: GlaS  $\rightarrow$  CAMELYON16 for WSOL baselines with and without PixelCAM. The higher  $J$  is, the more classes are separated.

## Appendix E. More Localization Performance Analysis

PxAP is the primary metric used to measure localization performance. However, following (Rony et al., 2023), we include other pixel-wise performance measures that is true/false positives/negative rates. As shown previously, our method PixelCAM achieves better results in terms of PxAP performance. This improvement is observed in Tab. 5 with a higher number of true positive/negative rates compared to GradCAM++, LayerCAM, SAT, and NEGEV on the GlaS dataset. For the CAMELYON16 dataset, we observe that PixelCAM increases the true positive rate compared to standard methods while being competitive to NEGEV. However, we note a significant improvement when applying PixelCAM to the SAT method, with a considerable increase in true positives, thereby reducing the over-activation issue.

		GlaS				CAMELYON16			
Bottom-up WSOL	Venue	TP*	FN*	TN*	FP*	TP*	FN*	TN*	FP*
DeepMIL (Ilse et al., 2018) †	<i>icml</i>	63.4	36.6	76.3	23.6	<b>66.9</b>	<b>33.1</b>	89.6	10.4
DeepMIL w/ NEGEV (Belharbi et al., 2022a) ★	<i>midl</i>	<b>79.0</b>	<b>21.0</b>	75.5	24.5	<u>64.0</u>	<u>36.0</u>	<u>92.6</u>	<u>7.4</u>
DeepMIL w/ PixelCAM †	-	<u>76.4</u>	<u>23.6</u>	<b>77.4</b>	<b>22.6</b>	59.0	41.0	<b>93.3</b>	<b>6.7</b>
GradCAM++ (Chattopadhyay et al., 2018) †	<i>wacv</i>	<u>62.0</u>	<u>38.0</u>	<b>79.9</b>	<b>20.1</b>	42.1	57.8	89.4	10.6
GradCAM w/ NEGEV (Belharbi et al., 2022a) ★	<i>midl</i>	58.4	41.6	76.0	24.0	<b>53.5</b>	<b>46.5</b>	<b>92.2</b>	<b>7.8</b>
GradCAM w/ PixelCAM †	-	<b>76.9</b>	<b>23.1</b>	<u>78.8</u>	<u>21.2</u>	<u>49.2</u>	<u>50.8</u>	<u>91.4</u>	<u>8.6</u>
LayerCAM (Jiang et al., 2021) †	<i>ieee tip</i>	62.3	37.7	<b>72.6</b>	<b>27.4</b>	29.6	70.3	86.8	13.2
LayerCAM w/ NEGEV (Belharbi et al., 2022a) ★	<i>midl</i>	<u>66.3</u>	<u>33.7</u>	70.7	29.3	<u>54.4</u>	<u>45.6</u>	<b>90.9</b>	<b>9.1</b>
LayerCAM w/ PixelCAM †	-	<b>77.6</b>	<b>22.4</b>	<u>72.5</u>	<u>27.5</u>	<b>56.5</b>	<b>43.5</b>	<u>89.2</u>	<u>10.8</u>
SAT (Wu et al., 2023) †	<i>iccv</i>	<u>58.8</u>	<u>41.2</u>	<u>63.8</u>	<u>36.2</u>	<u>20.0</u>	<u>80.0</u>	<b>92.2</b>	<b>7.8</b>
SAT w/ PixelCAM †	-	<b>67.6</b>	<b>32.4</b>	<b>75.0</b>	<b>25.0</b>	<b>65.7</b>	<b>34.2</b>	<u>79.2</u>	<u>20.8</u>
U-Net (Ronneberger et al., 2015)	<i>miccai</i>	88.9	11.1	89.8	10.2	68.0	32.0	94.5	5.5

Table 5: Confusion matrix performance over GlaS and CAMELYON16 test set with standard WSOL setup. † refers to model with a single-step approach while ★ refers to model from two-step family.

## Appendix F. Ablations

We provide in this section several ablations of our method `PixelCAM`.

### F.1. Impact of `PixelCAM` Depth

All the reported results of our method are obtained with a linear classifier that classifies pixel-embeddings into FG/BG classes. In this section, we further explore the impact of using a multi-layer classifier composed of three  $1 \times 1$  convolutional layers, which act as fully connected layers by reducing the dimensionality of the previous layer’s output by a factor of 2.

WSOL models	Venue	GlaS		CAMELYON16	
		PxAP	CL	PxAP	CL
DeepMIL (Ilse et al., 2018)	<i>icml</i>	79.9	<b>100.0</b>	71.3	85.0
DeepMIL w/ linear <code>PixelCAM</code>	-	<u>85.5</u>	<b>100.0</b>	<b>75.7</b>	<u>88.2</u>
DeepMIL w/ multi-layer <code>PixelCAM</code>	-	<b>86.3</b>	95.0	74.4	<b>88.6</b>
GradCAM++ (Chattopadhyay et al., 2018)	<i>wacv</i>	77.9	<b>100.0</b>	49.1	63.4
GradCAM++ w/ linear <code>PixelCAM</code>	-	<b>86.6</b>	<b>100.0</b>	<u>63.4</u>	<b>88.7</b>
GradCAM++ w/ multi-layer <code>PixelCAM</code>	-	86.5	95.0	<b>64.0</b>	<u>85.8</u>

Table 6: Localization (PxAP) and classification (CL) accuracy on GlaS and CAMELYON16 test set with linear and multi-layer `PixelCAM` for Standard WSOL setup.

As shown in Tab. 6, adding hidden layers to the pixel classifier is not necessarily an advantage for the localization task. Indeed, models with multiple layers perform worse on GlaS and CAMELYON16 for the WSOL baseline methods GradCAM++ (Chattopadhyay et al., 2018) and DeepMIL (Ilse et al., 2018), respectively. In the case of OOD data, Tab. 7 shows that large performance degradation can be observed when using multi-layer classifier. Therefore, as a general rule, we recommend using simply a linear pixel-classifier.

WSOL models	Venue	CAMELYON16 $\rightarrow$ GlaS		GlaS $\rightarrow$ CAMELYON16	
		PxAP	CL	PxAP	CL
DeepMIL (Ilse et al., 2018)	<i>icml</i>	64.5	<u>81.2</u>	29.0	<b>55.2</b>
DeepMIL w/ linear <code>PixelCAM</code>	-	<b>69.1</b>	<b>83.8</b>	<b>30.2</b>	<u>52.5</u>
DeepMIL w/ multi-layer <code>PixelCAM</code>	-	<u>67.4</u>	80.0	25.0	51.6
GradCAM++ (Chattopadhyay et al., 2018)	<i>wacv</i>	52.9	53.7	<b>39.1</b>	<u>52.4</u>
GradCAM++ w/ linear <code>PixelCAM</code>	-	<b>56.2</b>	<b>71.2</b>	<u>36.9</u>	<b>63.3</b>
GradCAM++ w/ multi-layer <code>PixelCAM</code>	-	<u>55.8</u>	<b>71.2</b>	33.8	50.6

Table 7: Localization (PxAP) and classification (CL) accuracy on GlaS and CAMELYON16 test set with linear and multi-layer `PixelCAM` for OOD setup.

### F.2. Impact of Model Selection of CAM Pre-trained Model

Our method leverages pseudo-labels extracted from a pretrained WSOL CAM-based model. However, due to the asynchronous convergence of classification and localization tasks (Rony et al., 2023), the criterion used for model selection is expected to have an impact on CAM quality, and therefore, pseudo-labels accuracy. Typically, models are selected based either on their classification performance on validation set by taking the best classifier (B-CL), or their localization performance and considering the best localizer (B-LOC). In Tab.8, we present the impact of such choice on our method PixelCAM.

WSOL models	Venue	GlaS		CAMELYON16	
		PxAP	CL	PxAP	CL
LayerCAM (Jiang et al., 2021)	<i>ieee tip</i>	75.1	<b>100.0</b>	33.2	84.8
LayerCAM w/ PixelCAM (B-LOC CAM)	-	<b>83.6</b>	<b>100.0</b>	<u>66.2</u>	<u>89.1</u>
LayerCAM w/ PixelCAM (B-CL CAM)	-	<u>80.6</u>	98.8	<b>67.5</b>	<b>89.4</b>
SAT (Wu et al., 2023)	<i>iccv</i>	65.9	98.8	32.8	83.2
SAT w/ PixelCAM (B-LOC CAM)	-	<b>79.1</b>	<b>100.0</b>	<b>51.2</b>	<b>87.2</b>
SAT w/ PixelCAM (B-CL CAM)	-	<u>75.5</u>	<b>100.0</b>	<u>35.6</u>	<u>86.4</u>

Table 8: Impact of model selection of CAM-based model to build pseudo-labels: Comparison of PixelCAM on GlaS and CAMELYON16 datasets on test set by using CAMs from B-LOC and B-CL models respectively. The performance is compared to the standard WSOL setup.

Initially, the CAMs used for pixel selection were generated from the B-LOC model. We compared the performance compared to using B-CL’s CAMs. Table 8 shows the results of various experiments with different WSOL baseline methods. We noticed that PixelCAM improves localization performance regardless of whether B-LOC or B-CL is used. However, for techniques that use gradients, such as GradCAM++ (Chattopadhyay et al., 2018) and LayerCAM (Jiang et al., 2021), CAMs generated with B-CL provide better performance. This can be explained by the nature of the CAMs produced by B-LOC for these methods. Specifically, B-LOC generates CAMs with better localization performance due to higher true positive rates but also introduces more false positives. This negatively impacts training, as incorrect labels are incorporated, potentially reducing overall performance.

We also observe a significant difference for SAT method (Wu et al., 2023). This is mainly explained by the fact that the CAMs generated by the B-CL model produce a high number of false positives, similar to those generated by B-LOC. However, the B-LOC model generates a higher true positive rate than B-CL, which is crucial for the strong performance of PixelCAM.

### F.3. Impact of Pixel Sampling Technique

Our method PixelCAM uses pseudo-annotation for the pixel classifier’s training. In particular, we consider random sampling of pixel locations to generate pseudo-labels as it has shown to yield better performance than fitting static regions (Belharbi et al., 2022a). This avoids

overfitting regions and promote exploring potential ROIs. In this section, we investigate the impact of various sampling techniques to identify pixels associated to FG and BG. We consider two different sampling approaches namely *Threshold-based* (Belharbi et al., 2022c) and *Probability-based (PB)* (Belharbi et al., 2022a). The *Threshold-based (TH)* approach automatically thresholds the CAM. Then, it considers all the pixels with activation above the threshold as foreground and valid for FG sampling. This is delineated by a mask. FG pixels are sampled uniformly within that mask, while BG pixels are sampled from outside the mask. The *Probability-based* approach samples FG pixels proportionally to the probability values obtained from the CAM using a multinomial distribution. However, the BG pixels are samples from  $(1 - \text{CAM})$  activations so regions with low activations will have higher sampling chance.

WSOL models	Venue	GlaS		CAMELYON16	
		PxAP	CL	PxAP	CL
GradCAM++ (Chattopadhyay et al., 2018)	<i>wacv</i>	76.8	<b>100.0</b>	49.1	63.4
GradCAM++ w/ TH PixelCAM	-	85.4	<b>100.0</b>	<u>54.2</u>	<b>89.9</b>
GradCAM++ w/ PB PixelCAM	-	<b>86.6</b>	<b>100.0</b>	<b>64.1</b>	<u>85.1</u>
SAT (Wu et al., 2023)	<i>iccv</i>	65.9	<u>98.8</u>	32.8	<u>83.2</u>
SAT w/ TH PixelCAM	-	<u>71.5</u>	92.5	<u>50.9</u>	81.0
SAT w/ PB PixelCAM	-	<b>79.1</b>	<b>100.0</b>	<b>51.2</b>	<b>87.2</b>

Table 9: Impact of pixel sampling technique on PixelCAM performance: Threshold-based (TH) vs. probability-based (PB), over standard WSOL setup.

In Tab. 9, we observe that pixel sampling technique has a significant impact on the localization performance of PixelCAM. For GlaS dataset, using GradCAM++ CAMs has a small impact, but for CAMs generated from a transformer architecture as SAT the impact is significant with a difference of 7.6%. Similarly, on a more challenging dataset as CAMELYON16, using a probabilistic approach for sampling pixels can lead to an improvement of 9.9%. This can be explained by the fact that sampling without relying on a threshold favors relevant pixels and most likely to have the correct pseudo-label. On the other hand, threshold-based method fixes a region for sampling with high likelihood of covering wrong regions. Learning with incorrect labels leads to poor models, and, therefore poor performance.

#### F.4. Impact of the Number of Sampled Pixels $n$

We analyze the impact of the number of pixels selected as pseudo-labels to train PixelCAM. To avoid unbalanced pixel classification, we sample the same number  $n$  of pixels as FG and BG. We consider the following cases  $n \in \{1, 5, 10, 20\}$ . As observed in Tab. 10, increasing  $n$  can affect both localization and classification with different degree, and depending on the dataset. In terms of localization, both datasets can slightly be affected. However, CAMELYON16 is largely affected in term of classification as performance can vary between 85.7% and 90.1%.



WSOL models	Venue	GlaS		CAMELYON16	
		PxAP	CL	PxAP	CL
GradCAM++ (Chattopadhyay et al., 2018)	<i>wacv</i>	76.8	<b>100.0</b>	49.1	63.4
GradCAM++ w/ PixelCAM:	-				
$n = 1$		<b>86.7</b>	<b>100.0</b>	63.5	<b>90.1</b>
$n = 5$		<u>86.6</u>	<b>100.0</b>	<u>64.1</u>	85.1
$n = 10$		86.0	<b>100.0</b>	63.2	<u>88.8</u>
$n = 20$		86.0	<b>100.0</b>	<b>64.7</b>	88.7

Table 10: Impact of the number of sampled pixel as pseudo-labels  $n$  on PixelCAM performance. We measure the PxAP and CL performance over the test set for GlaS and CAMELYON16 for standard WSOL setup.

### F.5. Impact of $\lambda$

Table 11 shows the impact of  $\lambda$  on the performance of our method PixelCAM. It achieves better results when using a higher value of  $\lambda$  on GlaS. As  $\lambda$  decreases, the PxAP performance also declines. On the CAMELYON16 dataset, we note that using a low  $\lambda$  value (0.001) significantly impacts localization performance. Therefore, we recommend using  $\lambda$  values in the range of 0.1 to 1.0 for better performance.

$\lambda$	GlaS		CAMELYON16	
	PxAP	CL	PxAP	CL
1	<b>83.6</b>	<b>100.0</b>	66.2	89.1
0.5	<u>83.1</u>	<b>100.0</b>	66.6	<b>89.7</b>
0.1	82.2	<b>100.0</b>	<b>67.4</b>	88.3
0.01	79.3	96.3	<u>66.7</u>	89.1
0.001	76.1	98.8	61.6	<u>89.2</u>

Table 11: Impact of hyper-parameter  $\lambda$  over PixelCAM in terms of PxAP and CL performance over test set. PixelCAM uses LayerCAM (Jiang et al., 2021) for pseudo-labels.

## Appendix G. More Visualization

### G.1. Visualization Results on Standard WSOL setup

In this section, we present visual results of **PixelCAM** compared to WSOL baseline methods. In a standard WSOL setup, we observe that **PixelCAM** enhances the ROIs detected by WSOL baseline methods on the **GlaS** dataset (Fig: 7 and 8). Regarding **CAMELYON16**, a challenging dataset, **PixelCAM** improves localization for cancer images by extending ROIs and reducing incorrect predictions from WSOL baseline methods as observed in (Fig: 9).




Figure 7: Standard WSOL setup. First column: **Normal** images from **GlaS**. Second column: Ground truth. Next columns: We display the visual CAM results for WSOL baseline without and with **PixelCAM**, respectively.






Figure 8: Standard WSOL setup. First column: **Cancerous** images from GlaS. Second column: Ground truth. Next columns: We display the visual CAM results for WSOL baseline without and with PixelCAM, respectively.




Figure 9: Standard WSOL setup. First column: **Normal** images from CAMELYON16. Second column: Ground truth. Next columns: We display the visual CAM results for WSOL baseline without and with PixelCAM, respectively.



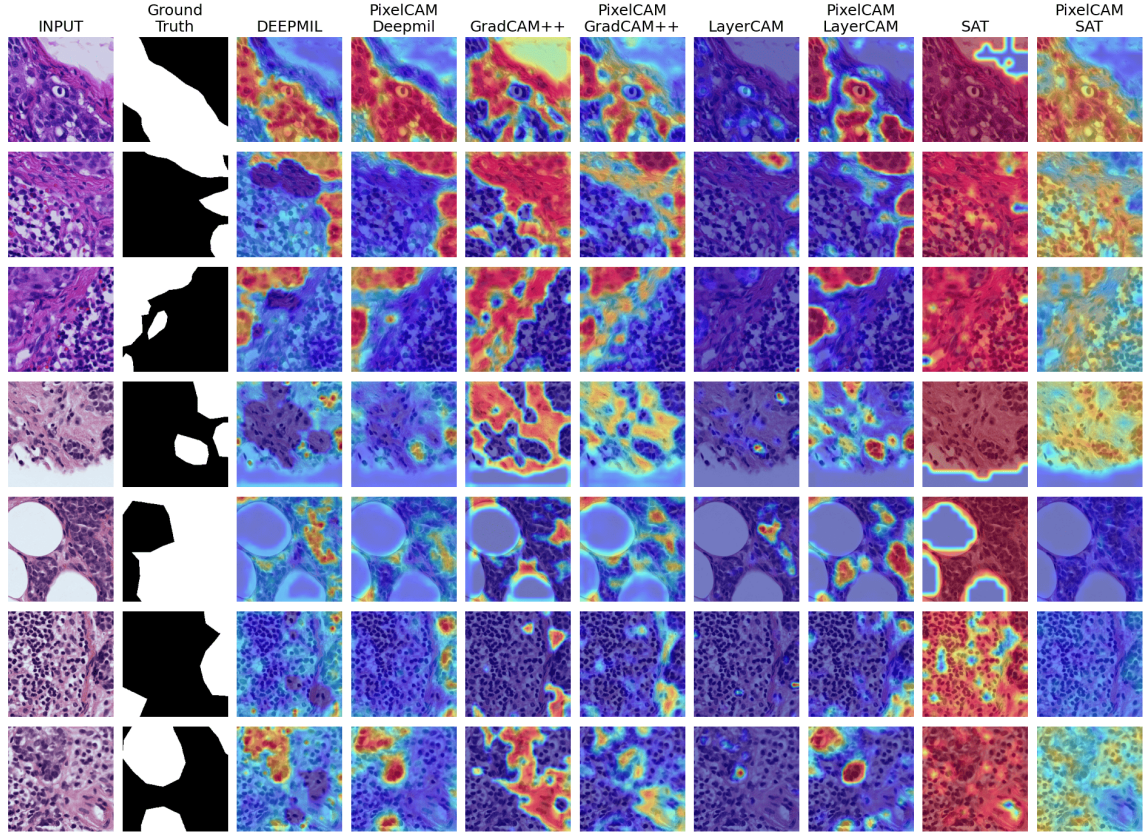


Figure 10: Standard WSOL setup. First column: **Cancerous** images from CAMELYON16. Second column: Ground truth. Next columns: We display the visual CAM results for WSOL baseline without and with PixelCAM, respectively.

## G.2. Visualization Results on Target set with OOD setup

We provide visual results on target set in the case of OOD for both scenarios: CAMELYON16  $\rightarrow$  GLaS and GLaS  $\rightarrow$  CAMELYON16. As we observe, PixelCAM can predict correctly in average cancer ROIs as illustrated in Fig: 12 and 14 but struggle with normal images as the WSOL baseline methods (Fig: 11 and 13).




Figure 11: OOD setup: CAMELYON16  $\rightarrow$  GLaS. First column: **Normal** images from GLaS. Second column: Ground truth. Next columns: We display the visual CAM results for WSOL baseline without and with PixelCAM, respectively.






Figure 12: OOD setup: GLaS  $\rightarrow$  CAMELYON16. First column: **Cancerous** images from GLaS. Second column: Ground truth. Next columns: We display the visual CAM results for WSOL baseline without and with PixelCAM, respectively.




Figure 13: OOD setup:  $\text{GlaS} \rightarrow \text{CAMELYON16}$ . First column: **Normal** images from CAMELYON16. Second column: Ground truth. Next columns: We display the visual CAM results for WSOL baseline without and with PixelCAM, respectively.



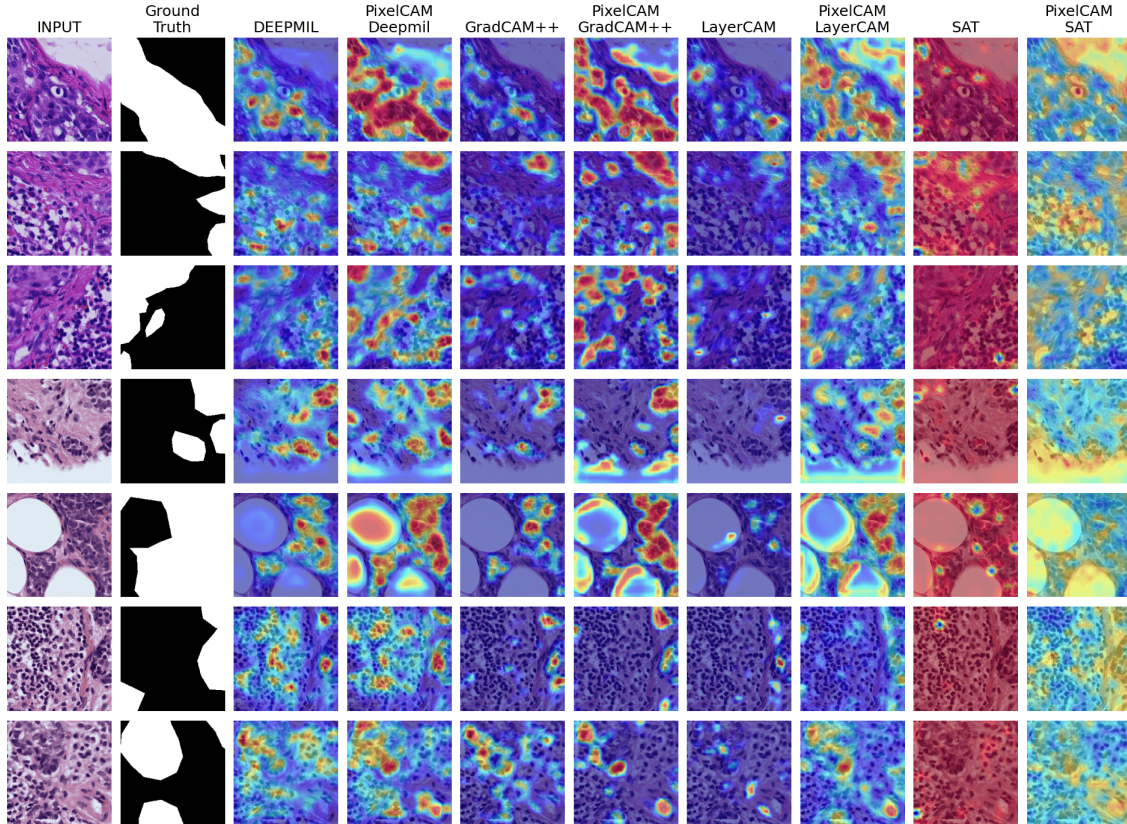


Figure 14: OOD setup: GLaS  $\rightarrow$  CAMELYON16. First column: **Cancerous** images from CAMELYON16. Second column: Ground truth. Next columns: We display the visual CAM results for WSOL baseline without and with PixelCAM, respectively.

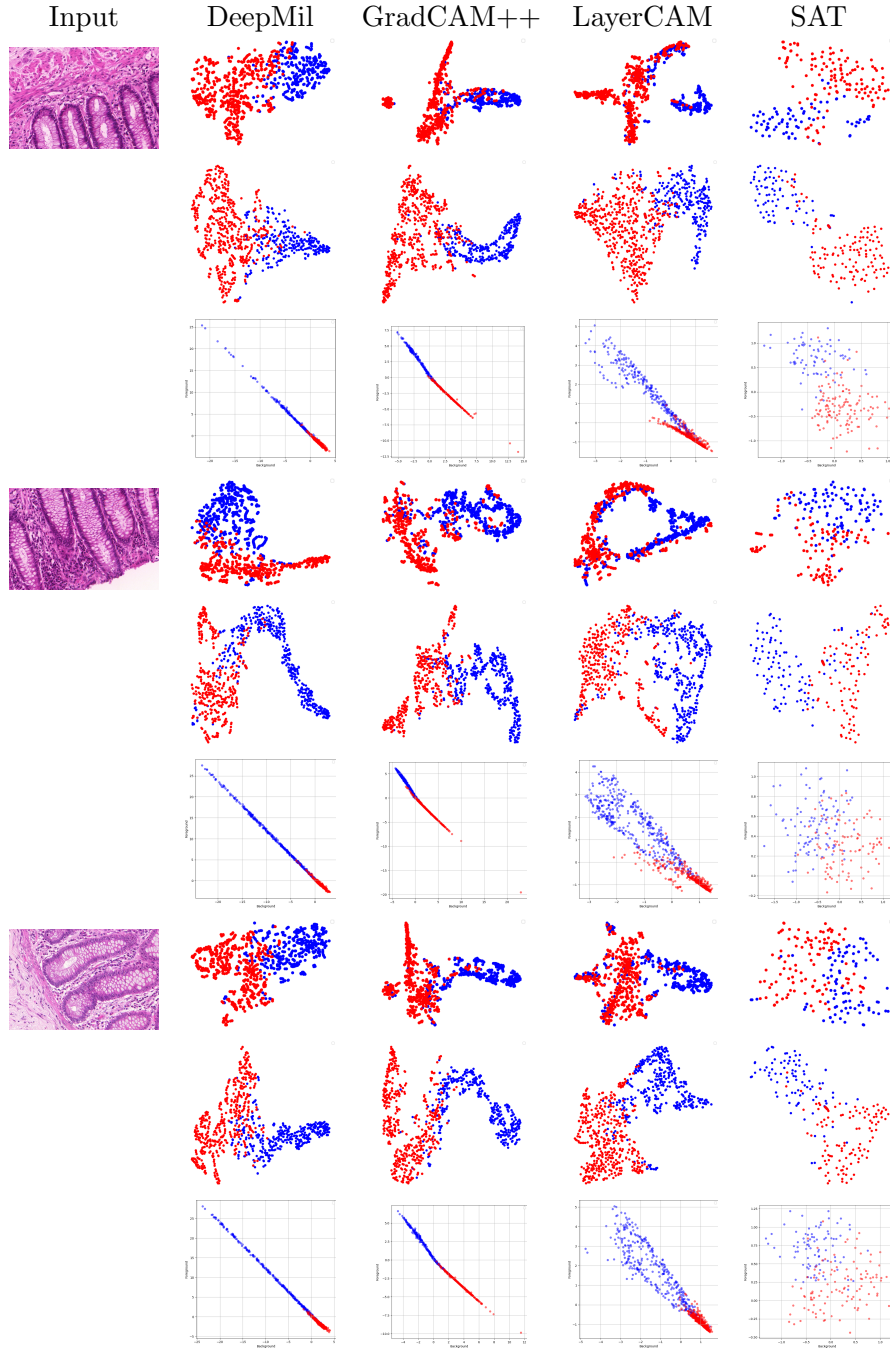


Figure 15: Standard WSOL setup: For the three **normal** images from GLaS, we display the t-SNE projection of foreground and background pixel-features of WSOL baseline methods without (1st row) and with (2nd row) PixelCAM. The 3rd row presents the logits of the pixel classifier of PixelCAM.

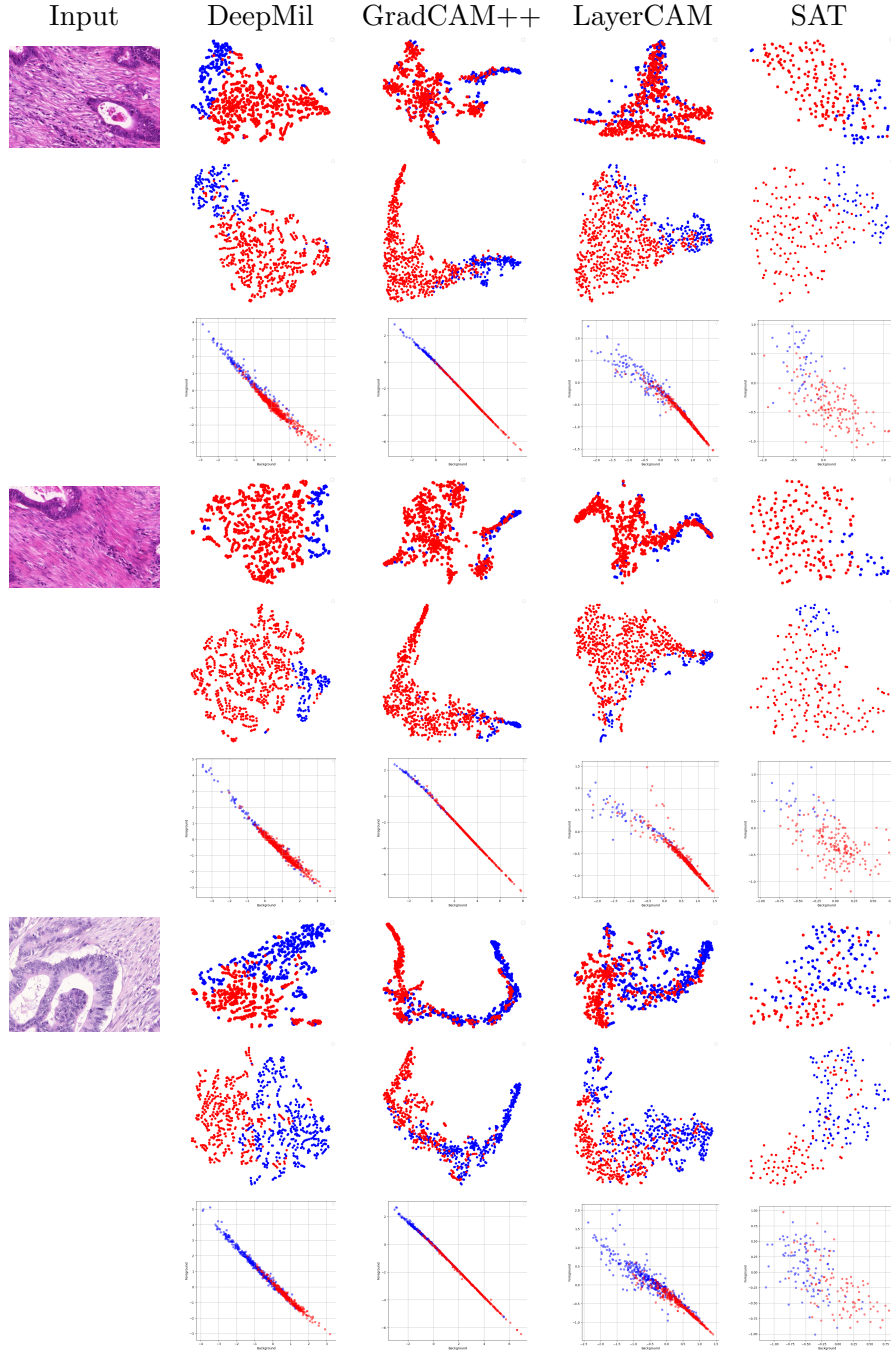


Figure 16: Standard WSOL setup: For the three **cancerous** images from **GlaS**, we display the t-SNE projection of foreground and background pixel-features of WSOL baseline methods without (1st row) and with (2nd row) **PixelCAM**. The 3rd row presents the logits of the pixel classifier of **PixelCAM**.

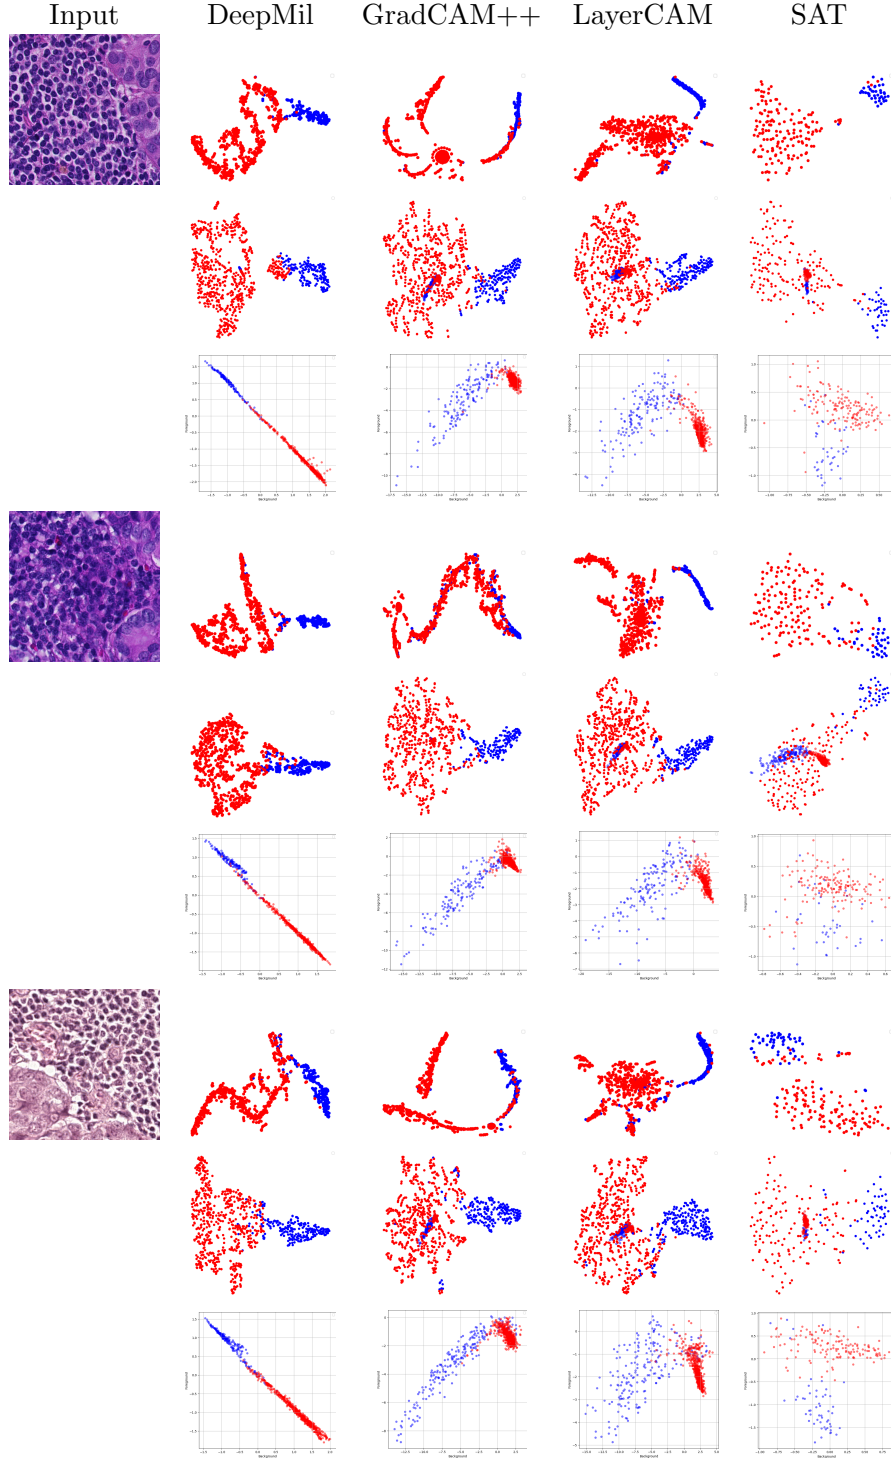


Figure 17: Standard WSOL setup: For the three **cancerous** images from CAMELYON16, we display the t-SNE projection of foreground and background pixel-features of WSOL baseline methods without (1st row) and with (2nd row) PixelCAM. The 3rd row presents the logits of the pixel classifier of PixelCAM.

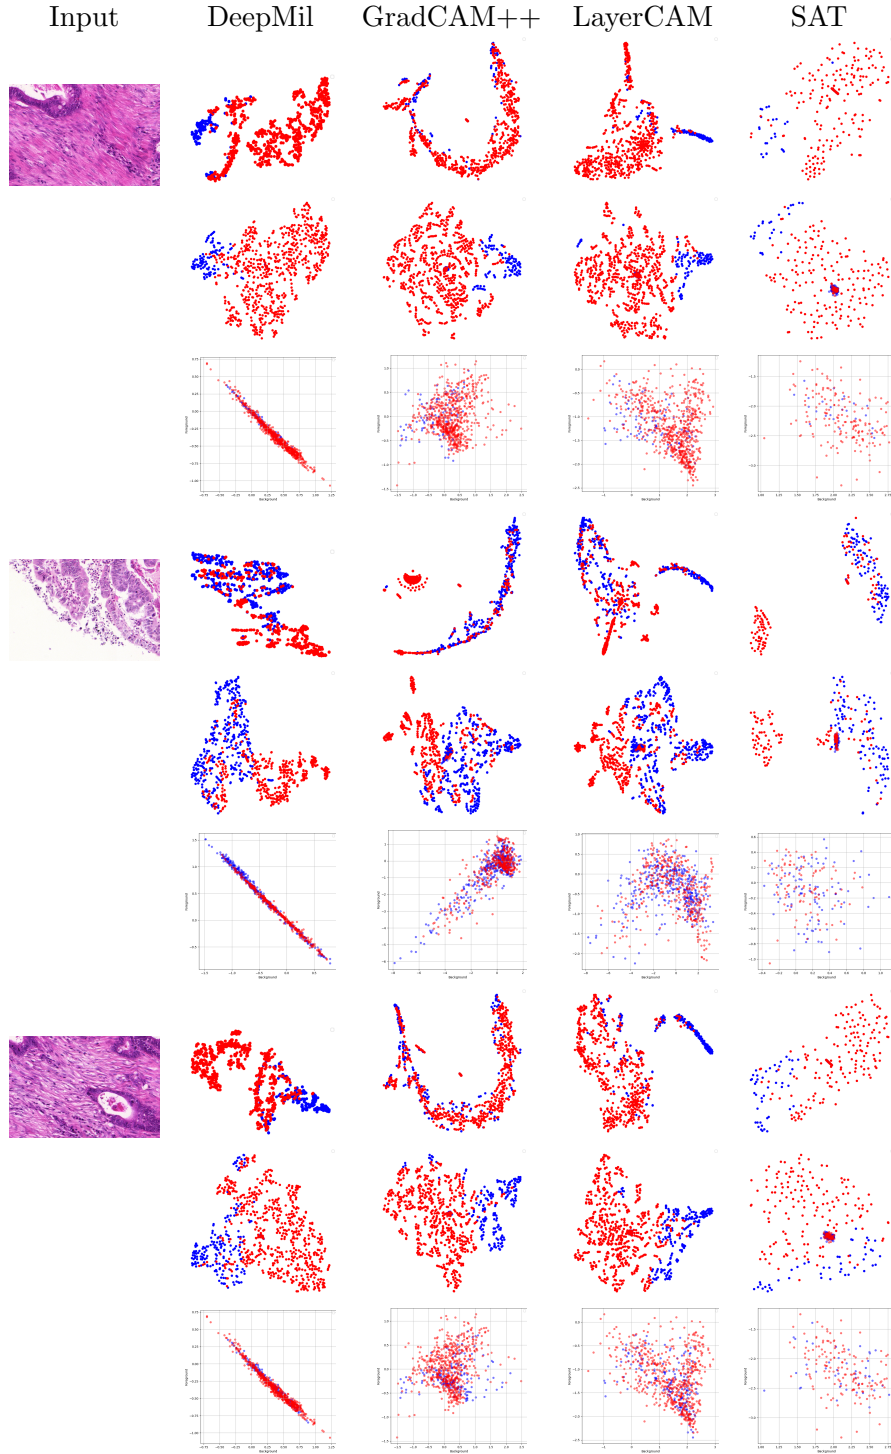


Figure 18: OOD setup (CAMELYON16  $\rightarrow$  GlaS): For the three **cancerous** images from GlaS, we display the t-SNE projection of foreground and background pixel-features of WSOL baseline methods without (1st row) and with (2nd row) PixelCAM. The 3rd row presents the logits of the pixel classifier of PixelCAM.



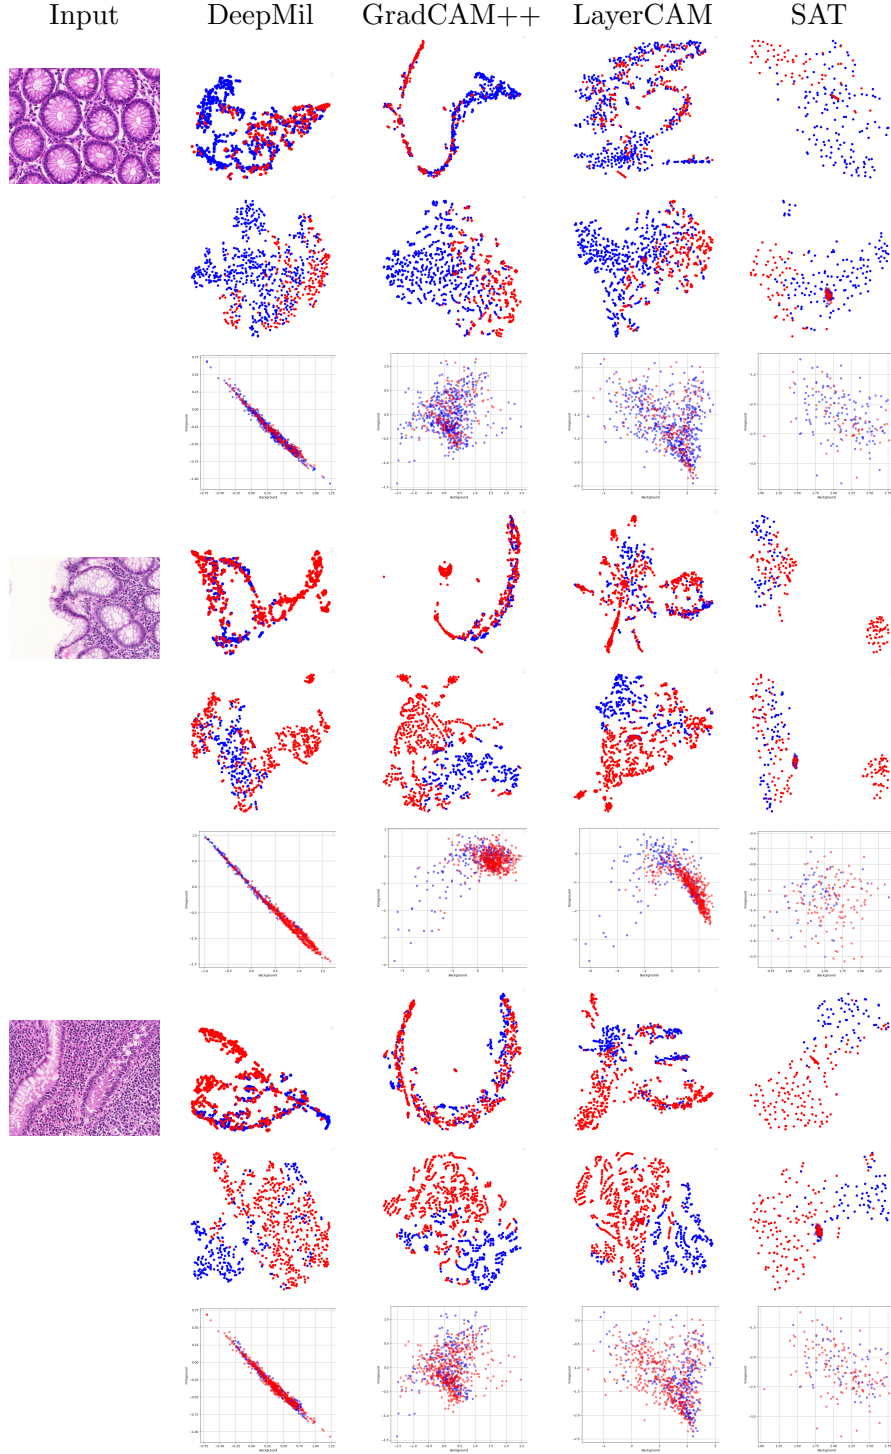


Figure 19: OOD setup (CAMELYON16  $\rightarrow$  GlaS): For the three **normal** images from GlaS, we display the t-SNE projection of foreground and background pixel-features of WSOL baseline methods without (1st row) and with (2nd row) PixelCAM. The 3rd row presents the logits of the pixel classifier of PixelCAM.

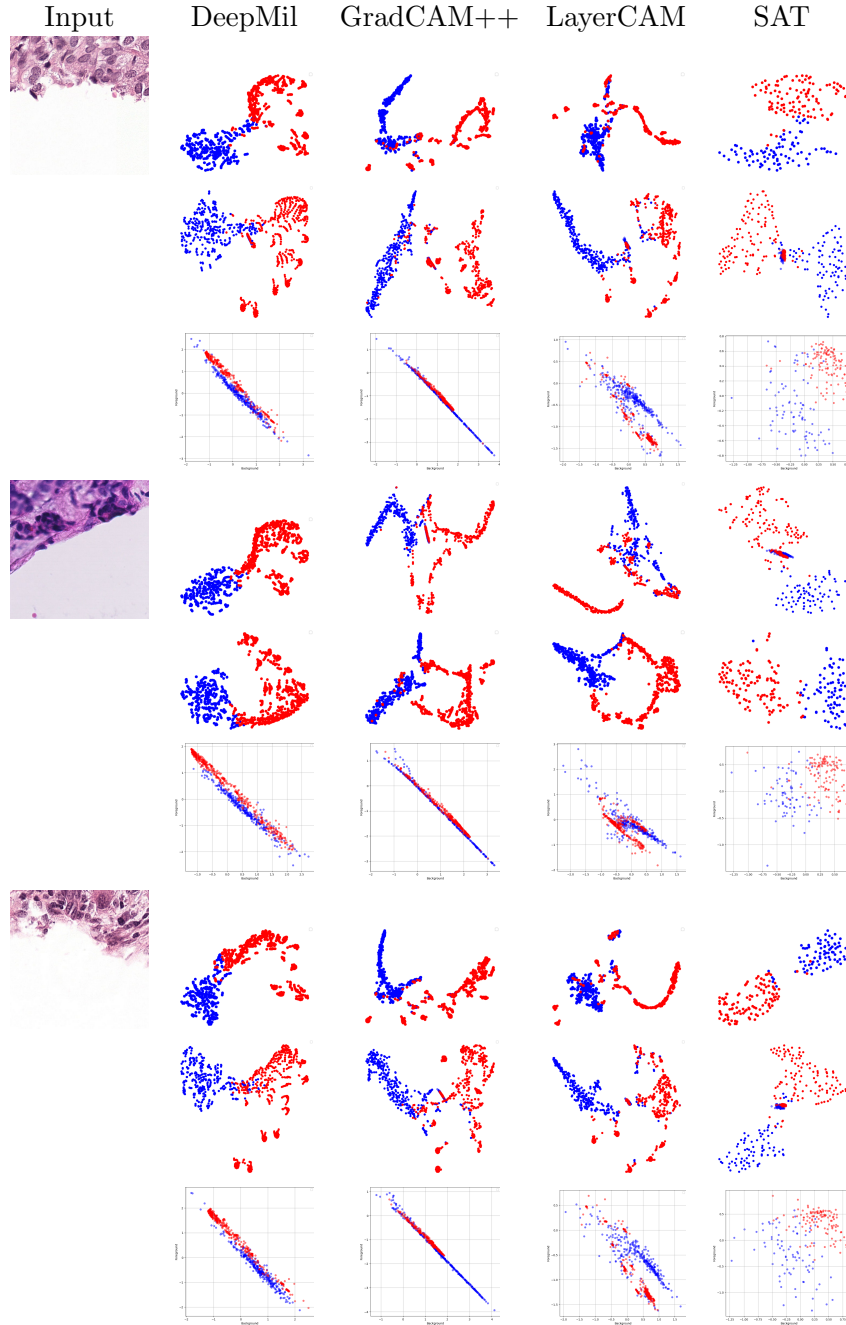


Figure 20: OOD setup (GlaS  $\rightarrow$  CAMELYON16): For the three **cancerous** images from CAMELYON16, we display the t-SNE projection of foreground and background pixel-features of WSOL baseline methods without (1st row) and with (2nd row) PixelCAM. The 3rd row presents the logits of the pixel classifier of PixelCAM.