

# ProMIL: Probabilistic Multiple Instance Learning for Medical Imaging

Łukasz Struski<sup>ad,\*</sup>, Dawid Rymarczyk<sup>ae</sup>, Arkadiusz Lewicki<sup>bd</sup>, Robert Sabiniewicz<sup>cd</sup>, Jacek Tabor<sup>ad</sup> and Bartosz Zieliński<sup>a</sup>

<sup>a</sup>Faculty of Mathematics and Computer Science, Jagiellonian University

<sup>b</sup>Faculty of Applied Computer Science, University of Information Technology and Management in Rzeszow

<sup>c</sup>Department of Pediatric Cardiology and Congenital Heart Diseases, Medical University of Gdansk

<sup>d</sup>UES Ltd.

<sup>e</sup>Doctoral School of Exact and Natural Sciences, Jagiellonian University

ORCID ID: Łukasz Struski <https://orcid.org/0000-0003-4006-356X>,

Dawid Rymarczyk <https://orcid.org/0000-0002-8543-5200>,

Arkadiusz Lewicki <https://orcid.org/0000-0002-1922-1562>, Jacek Tabor <https://orcid.org/0000-0001-6652-7727>,

Bartosz Zieliński <https://orcid.org/0000-0002-3063-3621>

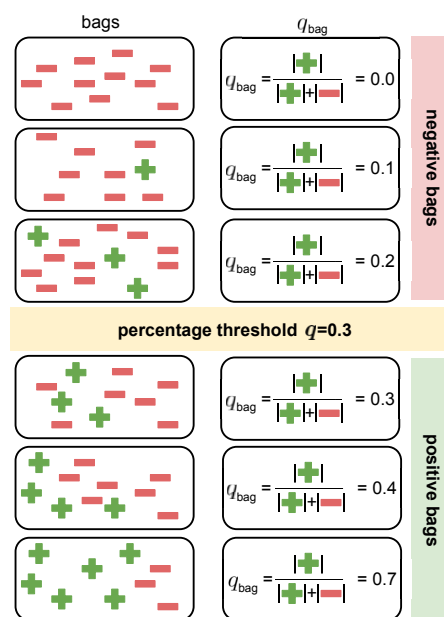
**Abstract.** Multiple Instance Learning (MIL) is a weakly-supervised problem in which one label is assigned to the whole bag of instances. An important class of MIL models is instance-based, where we first classify instances and then aggregate those predictions to obtain a bag label. The most common MIL model is when we consider a bag as positive if at least one of its instances has a positive label. However, this reasoning does not hold in many real-life scenarios, where the positive bag label is often a consequence of a certain percentage of positive instances. To address this issue, we introduce a dedicated instance-based method called ProMIL, based on deep neural networks and Bernstein polynomial estimation. An important advantage of ProMIL is that it can automatically detect the optimal percentage level for decision-making. We show that ProMIL outperforms standard instance-based MIL in real-world medical applications. We make the code available.

## 1 Introduction

Classification-oriented machine learning typically assumes that each example in a training set has a unique label assigned to it. However, in many practical situations, it is not feasible to label each instance individually. This results in a challenge known as Multiple Instance Learning (MIL) [9], where a bag of instances is associated with a single label, and it is assumed that some instances in the bag are relevant to that label.

MIL represents a powerful approach used in different application fields, mostly to solve problems where instances are naturally arranged in sets or to leverage weakly annotated data [4]. Its numerous applications include molecule classification task [9], predicting gene functions [11], content-based image retrieval [43], object location [17] and segmentation [31], computer-aided diagnosis and detection [2], document classification [14], and web mining [51].

Multitude applications of MIL contribute to introducing many new methods, described in the surveys [4, 50]. They generally divide into



**Figure 1:** In the percentage-based MIL assumption, each bag has a label indicating whether it is positive or negative based. It is based on the percentage of positive instances, but the information about the percentage threshold and the instance labels are unavailable during training. This example includes six various-sized bags containing positive and negative instances (green pluses and red minuses, respectively), and bags with a percentage of positive instances over threshold 0.3 are positive. In this paper, we define this assumption and introduce a dedicated method to solve it.

two types, instance- and representation-based. The former classifies instances and then aggregates those predictions to obtain a bag label. The latter aggregates instance representations and predicts a bag label directly. Therefore, instance-based approaches are more inter-

pretable but usually less accurate than representation-based methods. Most methods are representation-based and work only for the standard assumption, where a bag of instances is considered positive if it contains at least one positive instance. Only a few consider other assumptions described in [12], like presence-, threshold-, or count-based assumptions.

In this paper, we consider a specific type of MIL assumption, which we have named the percentage-based assumption. It is motivated by real-world medical problems, such as predicting neutrophils infiltration level in the Geboes scoring system on biopsy images [13], identifying bacteria strains on microscopic images [2], recording myocardial ischemia on ECG [41], or detecting congenital heart defects in newborns on ultrasound videos [21]. In percentage-based assumption, a bag is positive if a certain percentage of its instances is positive (however, this threshold is unknown during training). For instance, consider the evaluation of digestive tract health using the NHI scoring system [27], where a biopsy is given a score of 2 if more than 50% of crypts are infiltrated with neutrophils and there is no damage or ulceration in the epithelium.

The primary objective of this paper is to revisit the instance-based MIL methods and introduce a novel MIL assumption (percentage-based). We demonstrate the superiority of ProMIL over other instance-based methods, showcasing its excellent performance on benchmarks such as Doppler Ultrasound. These results highlight the importance of testing MIL methods beyond digital histopathology, as achieving high performance in medical imaging may necessitate different designs.

Furthermore, we argue that instance-based methods possess inherent interpretability since they provide a label for each instance. In contrast, ABMIL [18] only provides importance weights, and while extensions may provide an instance label, they may not be given directly (e.g. in DS-MIL [26]). With ProMIL, we aim to bridge this gap by directly indicating the required level of positive instances within a bag to classify it as positive.

We introduce a dedicated method ProMIL based on the Bernstein polynomial estimation that can be applied to any deep architecture. It classifies each instance separately and analyzes whether the appropriate percentage of those predictions corresponds to a positive label (is above the specific threshold). Importantly, ProMIL determines this threshold automatically during training.

We compared our approach with the baseline methods for artificial and real-world datasets, including our database with ultrasound videos of congenital heart defects in newborns that are among the most common developmental abnormalities and often lead to serious health consequences, including increased mortality and complications later in life. Finally, we particularly focus on the interpretable aspect of our method, bearing in mind its potential medical applications.

Our contributions can be summarized as follows:

- We define a novel percentage-based MIL assumption motivated by real-work medical problems.
- We introduce a method dedicated to this assumption based on the Bernstein polynomial estimator, which automatically discovers the threshold and can be applied to any deep architecture.
- Our method overpasses existing instance-based solutions, obtaining high performance while maintaining well-interpretable instance-level predictions.

## 2 Related Works

**Instance-based MIL.** Initial instance-based MIL methods trained a classifier by assigning pseudo labels (typically a bag label) to each instance. This approach was introduced in [30] and has been widely applied, e.g. in microscopy image classification [23] and drug activity prediction [49]. However, it can result in noisy labels. Therefore, recent studies focus on enhancing instance-based classifiers with dedicated MIL residual connections [44] or dynamic pooling [45]. Other approaches choose only crucial instances for training. Crucial instances can be selected based on classifier outputs [3], centered loss and centroids [6], or cluster conditioned distribution [34]. ProMIL is an instance-based MIL method, which, in contrast to existing solutions, is dedicated to percentage-based assumption and can automatically estimate the required threshold.

**Representation-based MIL.** Representation-based methods differ from instance-based approaches because they extract instance-level representations, aggregate them to a bag-level embedding, and use it to classify a bag directly. Most methods use the attention mechanism to assign importance to each instance during representation aggregation. AbMILP [18], introducing attention-based pooling, is one of the most popular approaches. It was extended by many mechanisms, including self-attention to detect intra-bag dependencies between instances [36], adversarial training to improve robustness [15], cluster assignments [52], siamese networks [46], attention loss [39], feature distillation [48], pyramidal fusion [26], prototypical parts for global interpretability [37, 47], multi-attention [22], and clustering constrained attention [29]. While there have been attempts to classify single instances while obtaining the bag-level embedding [32], the predictions obtained through this approach are still noisy. Hence, from the interpretability point of view, instance-based approaches, such as ProMIL, are preferable.

## 3 Preliminaries

**Multiple Instance Learning (MIL).** Multiple Instance Learning (MIL) is a weakly-supervised problem in which one label  $\mathbf{y} \in \mathbb{R}$  is assigned to a bag of instances  $\mathbf{x} = \{x_i\}_{i=0}^n$ , where  $n$  varies between bags [12]. Moreover, in the *standard MIL assumption*, the label of bag  $\mathbf{y} \in \{0, 1\}$ , each instance  $x_i$  has a hidden binary label  $y_i \in \{0, 1\}$  (unknown during training), and a bag is positive if at least one of its instances is positive

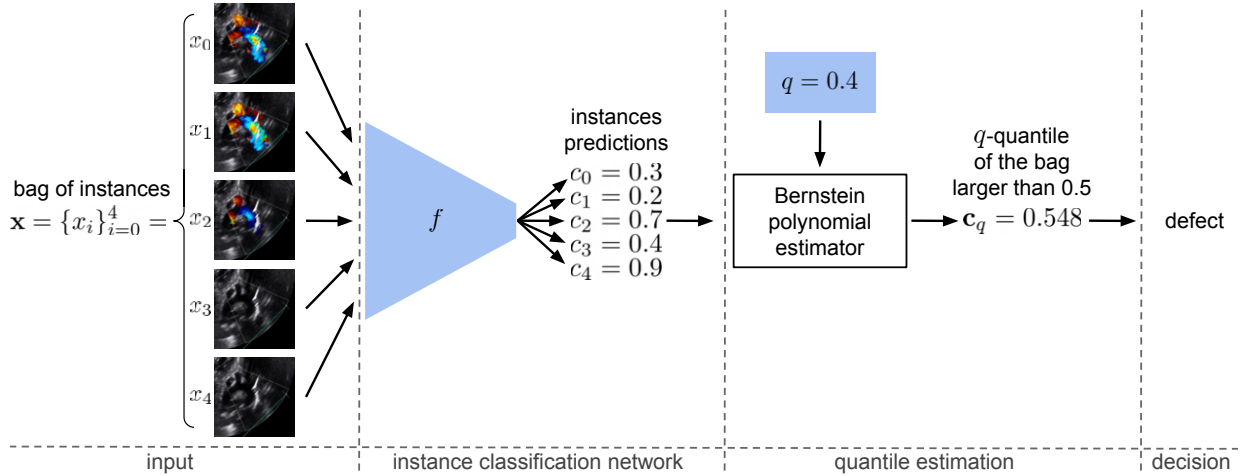
$$\mathbf{y} = \begin{cases} 0, & \text{iff } \sum_{i=0}^n y_i = 0, \\ 1, & \text{otherwise.} \end{cases} \quad (1)$$

**Bernstein polynomial estimator.** One of the main advantages of our approach is the ability to automatically discover the threshold in the percentage-based assumption. For this purpose, we use Bernstein polynomial estimator [5, 24, 53]. Assuming that we have a set  $\mathbf{p}$  with sorted elements  $p_0 \leq p_1 \leq \dots \leq p_n$ , its  $q$ -quantile estimator is calculated as

$$\mathbf{p}_q = \sum_{k=0}^n \binom{n}{k} q^{n-k} (1-q)^k \cdot p_k, \quad (2)$$

where  $\binom{n}{k}$  is a binomial coefficient.

Compared to the competitive solutions [19, 35, 42], Bernstein polynomial estimator provides more accurate estimates of quantiles



**Figure 2:** ProMIL model architecture consists of two trainable parameters (here marked as blue boxes), an instance classification network  $f$  and parameters  $q$ . The model takes a bag of instances (e.g. video frames) and passes them through network  $f$ , obtaining instances predictions  $c_0, c_1, \dots, c_4$ . These predictions are then used with the trainable parameter  $q$  in the differentiable estimation of the  $q$ -th quantile ( $c_q$ ). If the quantile exceeds 0.5, then we conclude that the bag is positive (e.g. correspond to a heart defect).

for datasets with complex or non-uniform distributions and is computationally efficient, even when implemented using standard numerical methods. These polynomials excel at capturing complex relationships between variables, making them well-suited for estimating quantiles, especially when the relationship between quantiles and the data distribution is non-linear. As the degree of the polynomial increases, the approximation of the true function improves, ensuring convergence. Moreover, Bernstein polynomials can be efficiently computed using de Casteljau’s algorithm [33], reducing computational complexity compared to other polynomial approximations.

In our numerically stable  $q$ -quantile estimator calculation (Eq. (2)), we leverage the properties of Bernstein polynomials. We represent all components of the equation as vectors. The logarithm of the  $\binom{n}{k}$  component is computed (see Eq. (4)). The two components,  $q^{n-k}$  and  $p_k$ , involve simple operations:  $(n - k) \cdot \log(q)$  and the *logsigmoid* value of the  $k$ -th instance, respectively. The term  $(1 - q)^k$  is calculated using the *log1mexp* function. Finally, since we represent all components of Eq. (2) as log-valued vectors, we utilize the *logsumexp* function to obtain the prediction value for a given bag (in logarithmic form)<sup>1</sup>.

## 4 ProMIL

Probabilistic Multiple Instance Learning (ProMIL) was created with the percentage-based assumption in mind that occurs in many real-world biomedical applications. It assumes that a positive bag is associated with a certain percentage of abnormal instances. Therefore, we classify those instances and then calculate the  $q$ -quantile of their predictions (in the range  $[0, 1]$ ). The bag is considered abnormal if the  $q$ -quantile is larger than 0.5 and normal otherwise. From this perspective, a method for quantile estimation should be differentiable to train the model and the optimal value of additional parameter  $q$ . Therefore, we use Bernstein polynomial estimation Eq. (2).

To formalize our approach, let us assume that  $\mathbf{x} = \{x_i\}_{i=0}^n$  is a bag of instances from the training set with the corresponding label  $\mathbf{y}$ . Moreover, let  $f$  be the classification network run for each instance  $x_i$  separately (see Figure 2).

---

### Algorithm 1: ProMIL training

---

**Data:**  $\mathbf{X}, \mathbf{Y}$  – bags from training set and their labels

**Parameter:**  $f$  – instance classification network;  $q$  – quantile (threshold)

**Result:** trained  $f$  and  $q$

- 1 randomly initialize  $q$ ;
  - 2 **while** no convergence **do**
  - 3   take random bag  $\mathbf{x} = \{x_0, x_1, \dots, x_n\}$  from  $\mathbf{X}$ ;
  - 4   take bag label  $\mathbf{y}$  from  $\mathbf{Y}$ ;
  - 5   calculate  $\mathbf{c} = \{c_i = f(x_i)\}_{i=0}^n$ ;
  - 6   sort  $\mathbf{c}$  in ascending order;
  - 7   calculate  $\mathbf{c}_q$  ( $q$ -quantile of  $\mathbf{c}$ ) using Eq. (2);
  - 8   calculate cost by applying  $\mathbf{c}_q$  and  $\mathbf{y}$  to Eq. (3);
  - 9   update  $f$  and  $q$  using cost back-propagation
- 

During training, presented in Algorithm 1, we take random bag  $\mathbf{x}$  from a training set and pass its instances through network  $f$ . As a result, we obtain a set of predictions  $\mathbf{c} = \{c_0, c_1, \dots, c_n\}$ . We sort this set so that  $c_0 \leq c_1 \leq \dots \leq c_n$ , and then we compute their  $q$ - and  $1 - q$ -quantile ( $\mathbf{c}_q$  and  $\mathbf{c}_{1-q}$ ) using Eq. (2). Because  $c_i \in [0, 1]$  then, according to properties of Bernstein polynomials [28]  $\mathbf{c}_q \in [0, 1]$ . Finally, we apply BCE loss to train our model

$$\text{cost}(\mathbf{x}, \mathbf{y}) = -\mathbf{y} \log \mathbf{c}_q - (1 - \mathbf{y}) \log \mathbf{c}_{1-q}, \quad (3)$$

where  $q$  is a trainable parameter.

Most of the calculations described in this algorithm are performed on a logarithmic scale for numerical stability, including *logsumexp* and the binomial coefficient calculated as


$$\binom{n}{k} = \exp(\log \Gamma(n+1) - \log \Gamma(k+1) - \log \Gamma(n-k+1)), \quad (4)$$

where  $\log \Gamma$  denotes the natural logarithm of the gamma function. Both *logsumexp* and *lgamma* functions are available in the PyTorch<sup>2</sup> framework.

To address a multi-class problem, Eq. (2) is defined for variables  $q$  and  $p_k$ , both represented in  $\mathbb{R}^m$ , where  $m$  corresponds to the number

<sup>1</sup> The source codes of our method are published on [github.com/gmum/ProMIL](https://github.com/gmum/ProMIL).

<sup>2</sup> <https://pytorch.org>



**Figure 3:** Four examples from the MNIST-bag dataset where the number of instances within a bag is sampled from  $\mathcal{N}(5, 2)$ , and the quantile  $q = 0.4$ . In red boxes, we highlight 9's corresponding to the positive instances.

of classes. These variables can be combined to form a vector  $\mathbf{p}_q$ , where each coordinate represents a specific class.

For the output of this model, Binary Cross-Entropy loss can be used. In the BCE scenario, each coordinate of  $\mathbf{p}_q$  is treated as an individual input to the BCE function. Notably, the use of vector processing techniques eliminates the need for explicit looping.

Furthermore, various operations mentioned in Algorithm 1, such as sorting and the computation of  $\mathbf{c}_q$ , can be extended to vectors or tensors with additional dimensions to accommodate multiple classes. This allows for streamlined calculations and facilitates the handling of class-specific information.


## 5 Experimental Setup

### 5.1 Datasets

**MNIST-bag** To evaluate the effectiveness of our ProMIL approach for percentage-based MIL assumption, we use the artificial MIL dataset MNIST-bag. The dataset consists of 1000 bags, each containing instances randomly sampled from a normal distribution with mean of 30 instances with variance equals to 5. The percentage of numbers of interests within a bag is taken randomly from a uniform distribution to ensure the balance within a dataset. We permute the order of instances within a bag to maintain permutation invariance. The examples of created bags from the dataset are presented in Figure 3.

The dataset is an enhancement of the dataset introduced in [18] with a different definition of a positive bag to fulfill the percentage-based assumption. We define the bag as positive if at least  $q$  fraction of instances from the bag are 9s. We maintain MNIST division into training and test sets.

**Histopathology datasets** Here we present details on three standard MIL datasets that represent MIL problems in real life, classification of histopathology images: Colon Cancer [40], Camelyon16 [10], and TCGA-NSCLC [1]. They consist of 100, 399, and 956 H&E stained whole slide images, respectively. For Colon Cancer, we use patches of the size  $32 \times 32$  while for other datasets we use  $224 \times 224$ . Colon Cancer has 100 bags with 2, 244 patches, Camelyon has 399 bags with a mean of 8, 871 patches, and TCGA-NSCLC has 1, 016



**Figure 4:** Ultrasound video frames corresponding to congenital heart defects in newborns, such as structural abnormalities and defects in large blood vessels.


bags with a mean of 3, 961 patches. From a medical perspective, the Colon Cancer dataset focuses on detecting epithelial nuclei which are often symptoms of developing disease, Camelyon16 is about the detection of micro- and macro-metastases in lymph nodes which is a prognostic factor for breast cancer growth. While the TCGA-NSCLC dataset is the identification of a cancer subtype which conditions which treatment should be applied. Each of those tasks is of high relevance from a clinical perspective.

**Doppler Ultrasound database** Our proposed approach holds significant potential in analyzing ultrasound videos (USG), which is currently one of the primary diagnostic tools for detecting congenital heart defects in newborns. These defects are among the most common developmental abnormalities and can lead to serious health consequences, including increased mortality and complications later in life.

The traditional approach to analyzing USG images involves individual examination and evaluation by qualified physicians specializing in pediatric cardiology. However, this method is time-consuming, requires highly specialized knowledge, and is susceptible to human errors. It is especially problematic for challenging cases, such as Tetralogy of Fallot, hypoplastic left heart syndrome, or defects associated with the pulmonary vein (see Figure 4).

We performed the Doppler ultrasound on over 250 patients to record blood flow in the heart and detect abnormalities. For some patients, the study was repeated, resulting in a dataset with over 1000 films of varying lengths (ranging from 10 frames to over 240 frames), with over 40% of the films being without defects. Heart defects specialists labeled each film. Before training the network, the data were randomly split into training and test sets while maintaining class proportions and ensuring that studies from the same patient were not included in both sets.

The training set constituted 85% of the entire dataset, while the remaining 15% was reserved for testing. Each frame was scaled to



**Figure 5:** The effectiveness of various MIL methods - AbMILP, Instance-MAX, Instance-MEAN, and ProMIL (our proposed method) on the MNIST-bag dataset was compared. The performance of these methods was analyzed under different scenarios, and the results were plotted. The left image shows the effect of the percentage of positive instances within a bag on the performance of the methods. The middle plot depicts how the bag size influences the effectiveness of MIL methods. The right plot shows the impact of the training dataset size on the model’s performance. It can be observed that our proposed ProMIL method consistently outperformed other methods in all scenarios. Moreover, AbMILP was found to be unstable due to overfitting when the number of instances within a bag increased. Lastly, the other methods showed higher variance in the results as compared to our ProMIL method.

200 × 200 pixels during network training, and a center crop of 172 pixels was used. During an evaluation, the frames were scaled to 172 × 172 pixels.

## 5.2 Implementation details

For the MNIST-bag dataset, we use LeNet-5 [25] network to encode the images. It is learned for 100 epochs with early stopping after 15 epochs without a change on a validation set. Training for each scenario was repeated five times. We use Adam optimizer [20] with learning rate  $10^{-4}$  with  $\beta_1 = 0.99$ ,  $\beta_2 = 0.999$ , weight decay equals  $10^{-5}$ , and batch size 1. AbMILP used as a baseline method also uses LeNet5 backbone, consisting of one attention head with embedding size equal to 500 and classifier hidden space 128. We use a binary cross entropy as a loss function.

For Colon Cancer dataset we use backbone from [40], Adam optimizer with learning rate  $10^{-3}$  with  $\beta_1 = 0.99$ ,  $\beta_2 = 0.999$ , weight decay equals  $10^{-3}$ , and batch size 1. We train a model for 200 epochs with an early stopping window equal to 25 epochs. We split the data into ten folds and repeat the training 5 times, similarly as it is done in [18].

When it comes to Camelyon16 and TCGA-NSCLC datasets, we use pretrained in a self-supervised way ResNet18 [7] as a backbone and train a model for 200 epochs with an early stopping window equal to 40 epochs. Also, we use Adam optimizer with learning rate  $5 \cdot 10^{-4}$  with  $\beta_1 = 0.99$ ,  $\beta_2 = 0.999$ , weight decay equals  $10^{-3}$ , and batch size 1.

For all the aforementioned datasets, we performed a hyperparameter random grid search with the following sets of parameters: learning rate in the range  $[10^{-6}, 10^{-2}]$ , weight decay in range  $[0, 1]$ ,  $q$  in the range  $[0.1, 0.5]$ .

To compute frame-based predictions for ultrasound videos, we utilized ResNet-18 [16], which was pretrained on the ImageNet [8] dataset. We replaced the last layer with two fully-connected layers, each having dimensions of 128 and 64, respectively. We employed grid search techniques to fine-tune hyperparameters, such as the learning rate (with possible values of  $10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}$ ),

dropout rate (searched over values of 0.1, 0.2, 0.3, 0.4), and L2 regularization (with possible values of 0,  $10^{-6}, 10^{-5}, 10^{-4}$ ). Moreover, we experimented with different learning rate schedulers, including CyclicLR, LambdaLR, and without any scheduler.

We trained the models in two stages. In the first stage, we froze the feature extraction part of the pre-trained ResNet-18 model and only trained the two additional linear layers for 25 epochs. In the second stage, we fine-tuned the entire network for 100 epochs. To train the models, we used the Adam optimizer [20] with default values, similar to other datasets. To accommodate the varying lengths of videos in our dataset, we utilized a batch size of 1.

## 6 Results

This section provides the results of experiments conducted on five different datasets. We begin with MNIST-bag, which is a toy dataset that enables a more thorough understanding and analysis of our ProMIL method. Next, we present our findings on three histopathological datasets, which are commonly used benchmarks for MIL algorithms, demonstrating the versatility of our approach. Finally, we introduce a novel heart ultrasound dataset, where we demonstrate the suitability of our ProMIL method for analyzing medical videos. We also include a detailed interpretability analysis of the method conducted with specialized medical doctors to confirm that our models accurately recognize important features.

**MNIST-bag results.** We conducted a comparative analysis of our ProMIL approach with the instance-max and instance-mean methods, as well as with Attention MIL Pooling [18]. As illustrated in Figure 5, our ProMIL outperforms all other methods in terms of AUC. Additionally, we found that percentage-based assumption is difficult for AbMILP and can lead to random performance on the test dataset, even if the model converges during training. Moreover, we observed that Instance-MAX and Instance-MEAN struggle to distinguish between positive and negative bags when the ground truth  $q$  is in the range of  $[0.1, 0.5]$ . With an increase in bag size or the number of training bags, all models achieve better performance. However, our



**Table 1:** Results for Colon Cancer dataset, where ProMIL outperforms the baseline methods (INSTANCE+MAX and INSTANCE+MEAN). However, it achieves worse results compared to more advanced techniques that employ attention mechanisms such as SA-ABMILP\*. Notice that values for comparison are indicated with “\*” from [37].

COLON CANCER		
METHOD	ACCURACY	AUC
EMBEDDING+MAX*	82.4% ± 1.5%	0.918 ± 0.010
EMBEDDING+MEAN*	86.0% ± 1.4%	0.940 ± 0.010
AbMILP*	88.4% ± 1.4%	0.973 ± 0.007
SA-AbMILP*	90.8% ± 1.3%	0.981 ± 0.007
PROTOMIL*	81.3% ± 1.9%	0.932 ± 0.014
INSTANCE+MAX*	84.2% ± 2.1%	0.914 ± 0.010
INSTANCE+MEAN*	77.2% ± 1.2%	0.866 ± 0.008
PROMIL (OUR)	<b>88.1% ± 1.7%</b>	<b>0.969 ± 0.010</b>

**Table 2:** Our ProMIL approach achieves comparable or better results than baseline approaches, specifically INSTANCE+MAX and INSTANCE+MEAN, for the Camelyon16 and TCGA-NSCLC datasets. However, it falls short of matching the performance of more complex models such as TransMIL, which are based on transformer architectures. It is worth noting that ProMIL matches the performance of the INSTANCE+MAX approach. This may be related to the characterization of lymph metastasis (Camelyon16 dataset), which typically covers a small tissue portion. As a result, the ProMIL learns to estimate a small value for  $q$  and achieves similar results to the INSTANCE+MAX approach. Notice that values for comparison marked with “\*” are taken from [37].

METHOD	CAMELYON16		TCGA-NSCLC	
	ACCURACY	AUC	ACCURACY	AUC
MILRNN*	80.62%	0.807	86.19%	0.910
AbMILP*	84.50%	0.865	77.19%	0.865
DSMIL*	86.82%	0.894	80.58%	0.892
CLAM-SB*	87.60%	0.881	81.80%	0.881
CLAM-MB*	83.72%	0.868	84.22%	0.937
TRANSMIL*	88.37%	0.931	88.35%	0.960
PROTOMIL*	87.29%	0.935	83.66%	0.918
INSTANCE+MEAN*	79.84%	0.762	72.82%	0.840
INSTANCE+MAX*	<b>82.95%</b>	<b>0.864</b>	85.93%	<b>0.946</b>
PROMIL (OUR)	<b>82.95%</b>	<b>0.864</b>	<b>86.71%</b>	0.940

**Table 3:** Our ProMIL method achieves state-of-the-art results for Doppler Ultrasound benchmark. Moreover, the lower than ProMIL performance of AbMILP and SA-AbMILP shows that our ProMIL is more suitable for this task. By ‘Accuracy’ we mean balanced accuracy.

DOPPLER ULTRASOUND		
METHOD	ACCURACY	AUC
AbMILP	79.77%	0.85
SA-AbMILP	84.85%	0.88
INSTANCE+MAX	70.09%	0.71
INSTANCE+MEAN	75.67%	0.79
PROMIL (OUR)	<b>86.36%</b>	<b>0.92</b>

ProMIL still outperforms them, making it suitable for both small and large datasets with small or large bags.

**Histopathological results.** We evaluated the versatility of our ProMIL approach on three real-life datasets commonly used as standard benchmarks in MIL: Colon Cancer, Camelyon16, and TCGA-NSCLC. The results for Colon Cancer are presented in Table 1, while those for Camelyon16 and TCGA-NSCLC are in Table 2. Our ProMIL approach outperformed other Instance-based methods for Colon Cancer and TCGA-NSCLC and matched the performance of the Instance-MAX approach for Camelyon16. This similarity in performance can be attributed to the small amount of metastatic tissue required for a positive label in Camelyon16 [10], leading our ProMIL to learn a small  $q$  what is making it similar to Instance-MAX. When compared to other MIL approaches that operate on the bag level, such as TransMIL and ProtoMIL, our ProMIL outperformed some of them (e.g., MILRNN for TCGA-NSCLC and AbMILP for Colon Cancer) and was slightly worse than recent approaches using transformer-based architectures such as TransMIL. These results demonstrate the adaptability of ProMIL to a variety of MIL problems in the digital pathology field.

**Doppler Ultrasound database results.** We conducted an evaluation of our approach on a dataset of ultrasound videos with varying numbers of frames. Similar to the previous experiment, we compared our method with the baseline (instance-based) approaches. The results of this evaluation are presented in Table 3.

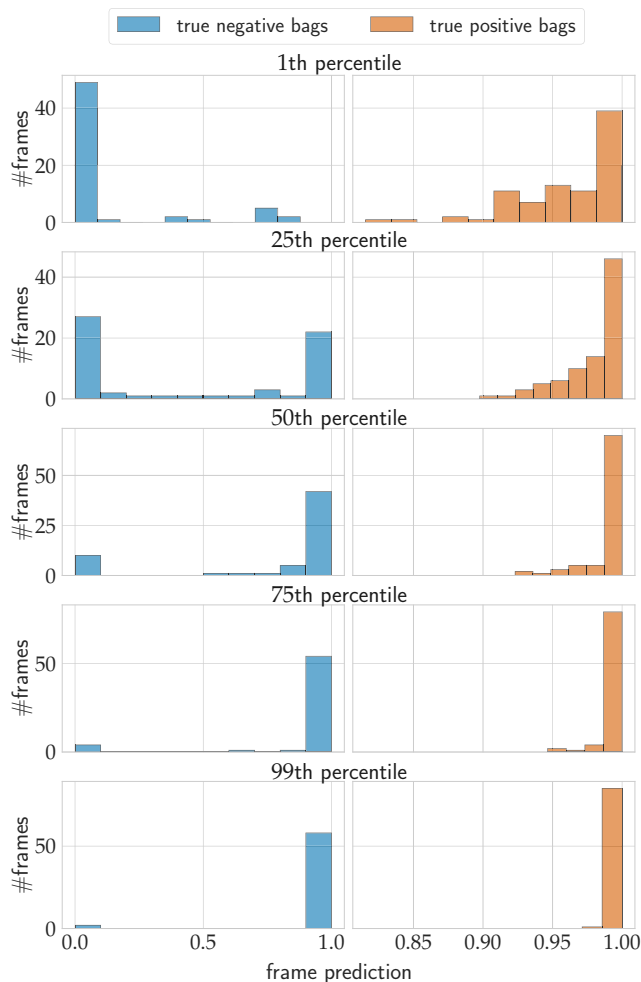
ProMIL utilizes a unique strategy for selecting the most significant frames in each video using the  $q = 38.56\%$ -quantile criterion. This approach contrasts with the competing methods, INSTANCE+MAX and INSTANCE+MEAN, which rely on single-frame and all-frame analyses, respectively. In our evaluation, we found that choosing the dominant prediction did not necessarily lead to high-performance scores. Furthermore, using all frames, as done in the INSTANCE+MEAN method, produced better results than the INSTANCE+MAX method, but still failed to achieve significantly better results, indicating overfitting to the noise.

Our approach’s superiority over the baseline methods is further highlighted by the effectiveness of selecting the most significant frames in each video, allowing us to mitigate the effects of noise and achieve superior performance.

Furthermore, Figure 6 provides insights into the distribution of the predictions made by the  $f$  classification component of the ProMIL model. This distribution is based on the selection of significant frames, which our model determined to be 38.56% for X data. To illustrate this, we only considered true negative and positive bags and presented their prediction distributions separately. Our results show that for true positive bags, all significant frames have predictions above 0.8, indicating high confidence in their classification. Conversely, for true negative bags, up to the 50th percentile, instances with predictions close to zero dominate the model’s inference.

In addition to analyzing the number and selection of frames, our study also focused on identifying the specific objects and regions of interest in individual frames that our model considers when making predictions. To achieve this, we utilized the widely recognized Grad-CAM technique [38]. This approach generates a heatmap highlighting the critical regions of an input image that the model focused on to arrive at its decision.

Figure 7 presents several results that demonstrate the heart areas our model emphasized when examining ultrasound images of a child with a congenital heart defect that involves a septal defect. After con-




**Figure 6:** The illustration depicts the classification predictions generated by the ProMIL  $f$  classification component. This component relies on the selection of significant frames, representing 38.56% of the frames with the highest prediction values, that impacted its decision-making process. The figure displays the 1st, 25th, 50th, 75th, and 99th percentiles of these significant frames in each row. The left column exhibits the prediction distribution for true negative bags, while the right column shows the distribution for true positive bags. Note that frames with predictions below threshold 0.5 are absent from the right column, which is not apparent from the left column. Even among the 99th percentile, there are frames with predictions approaching zero.

sulting with clinical experts in the field of pediatric cardiology, we confirmed that the regions marked on the heatmaps align with the areas typically identified by experts. Consequently, our approach could assist medical practitioners in comprehending the outcomes and recognizing heart defects in newborns.

## 7 Conclusions

In this paper, we propose a novel instance-based approach ProMIL to Multiple Instance Learning (MIL), where bag classification is obtained by aggregating instance-level predictions. It solves the percentage-based assumption, which occurs in many real-world medical applications, using Bernstein polynomial estimation.

ProMIL outperforms other instance-based approaches and obtains



**Figure 7:** An ultrasonographic image of a newborn’s heart with a ventricular septal defect and the corresponding heat map. This map allows to understand which features of the image were important for the model during the prediction process, taking into account the differences between successive frames of the processed and analyzed ultrasound recording. Areas with larger changes on the heat map are marked in yellow, red and burgundy as areas with higher heat intensity. Consultations with medical experts in the field of cardiology confirmed that such areas in the presented image indicate the locations of the ventricular septal defect.

results comparable to less interpretable representation-based models. Therefore, it can be easily applied to critical domains, such as computer-assisted interventions.

## 8 Acknowledgements

This research was partially funded by the National Science Centre, Poland, grants no. 2020/39/D/ST6/01332 (work by Łukasz Struski), 2022/47/B/ST6/03397 (work by Bartosz Zieliński), and 2022/45/N/ST6/04147 (work by Dawid Rymarczyk). Moreover, Dawid Rymarczyk received an incentive scholarship from the funds of the program Excellence Initiative – Research University at the Jagiellonian University in Kraków. Works of Jacek Tabor, Arkadiusz Lewicki, and Robert Sabiniewicz were supported by company the UES Ltd. Some experiments were performed on servers purchased with funds from a grant from the Priority Research Area (Artificial Intelligence Computing Center Core Facility) under the Strategic Programme Excellence Initiative at Jagiellonian University.

## References

- [1] Shaimaa Bakr, Olivier Gevaert, et al., ‘A radiogenomic dataset of non-small cell lung cancer’, *Scientific data*, **5**(1), 1–9, (2018).
- [2] Adriana Borowa, Dawid Rymarczyk, et al., ‘Identifying bacteria species on microscopic polyculture images using deep learning’, *IEEE Journal of Biomedical and Health Informatics*, **27**(1), 121–130, (2023).
- [3] Gabriele Campanella, Matthew G Hanna, et al., ‘Clinical-grade computational pathology using weakly supervised deep learning on whole slide images’, *Nature medicine*, **25**(8), 1301–1309, (2019).
- [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).
- [5] Cheng Cheng, ‘The bernstein polynomial estimator of a smooth quantile function’, *Statistics & probability letters*, **24**(4), 321–330, (1995).
- [6] Philip Chikontwe, Meejeong Kim, Soo Jeong Nam, Heounjeong Go, and Sang Hyun Park, ‘Multiple instance learning with center embeddings for histopathology classification’, in *MICCAI*, (2020).
- [7] Ozan Ciga, Anne L Martel, and Tony Xu, ‘Self supervised contrastive learning for digital histopathology’, *preprint arXiv:2011.13971*, (2020).
- [8] Jia Deng, Wei Dong, et al., ‘Imagenet: A large-scale hierarchical image database’, in *CVPR*, (2009).

- [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).
- [10] Babak Ehteshami Bejnordi, Mitko Veta, et al., 'Diagnostic Assessment of Deep Learning Algorithms for Detection of Lymph Node Metastases in Women With Breast Cancer', *JAMA*, **318**(22), 2199–2210, (12 2017).
- [11] Ridvan Eksi, Hong-Dong Li, et al., 'Systematically differentiating functions for alternatively spliced isoforms through integrating rna-seq data', *PLoS computational biology*, **9**(11), e1003314, (2013).
- [12] James Foulds and Eibe Frank, 'A review of multi-instance learning assumptions', *The knowledge engineering review*, **25**(1), 1–25, (2010).
- [13] Karel Geboes, R Riddell, A Öst, B Jensfelt, T Persson, and R Löfberg, 'A reproducible grading scale for histological assessment of inflammation in ulcerative colitis', *Gut*, **47**(3), 404–409, (2000).
- [14] Zellig S Harris, 'Distributional structure', *Word*, **10**(2-3), 146–162, (1954).
- [15] Noriaki Hashimoto, Daisuke Fukushima, et al., 'Multi-scale domain-adversarial multiple-instance cnn for cancer subtype classification with unannotated histopathological images', in *CVPR*, (2020).
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, 'Deep residual learning for image recognition', in *CVPR*, (2016).
- [17] Judy Hoffman, Deepak Pathak, Trevor Darrell, and Kate Saenko, 'Detector discovery in the wild: Joint multiple instance and representation learning', in *CVPR*, (2015).
- [18] Maximilian Ilse, Jakob Tomczak, and Max Welling, 'Attention-based deep multiple instance learning', in *ICML*, (2018).
- [19] Maurice George Kendall et al., 'The advanced theory of statistics.', *The advanced theory of statistics.*, (2nd Ed), (1946).
- [20] Diederik P Kingma and Jimmy Ba, 'Adam: A method for stochastic optimization', *preprint arXiv:1412.6980*, (2014).
- [21] Masaaki Komatsu, Akira Sakai, Reina Komatsu, et al., 'Detection of cardiac structural abnormalities in fetal ultrasound videos using deep learning', *Applied Sciences*, **11**(1), 371, (2021).
- [22] Andrei V Konstantinov and Lev V Utkin, 'Multi-attention multiple instance learning', *Neural Computing and Applications*, **34**(16), 14029–14051, (2022).
- [23] Oren Z Kraus, Jimmy Lei Ba, and Brendan J Frey, 'Classifying and segmenting microscopy images with deep multiple instance learning', *Bioinformatics*, **32**(12), i52–i59, (2016).
- [24] Alexandre Leblanc, 'On estimating distribution functions using bernstein polynomials', *Annals of the Institute of Statistical Mathematics*, **64**(5), 919–943, (2012).
- [25] Yann LeCun, Bernhard Boser, et al., 'Handwritten digit recognition with a back-propagation network', *NeurIPS*, (1989).
- [26] Bin Li, Yin Li, and Kevin W Eliceiri, 'Dual-stream multiple instance learning network for whole slide image classification with self-supervised contrastive learning', in *CVPR*, (2021).
- [27] Katherine Li, Richard Strauss, Colleen Marano, et al., 'A simplified definition of histologic improvement in ulcerative colitis and its association with disease outcomes up to 30 weeks from initiation of therapy: Post hoc analysis of three clinical trials', *Journal of Crohn's and Colitis*, **13**(8), 1025–1035, (2019).
- [28] G. G. Lorentz, Bernstein polynomials, 1953.
- [29] Ming Y Lu, Drew FK Williamson, et al., 'Data-efficient and weakly supervised computational pathology on whole-slide images', *Nature biomedical engineering*, **5**(6), 555–570, (2021).
- [30] Oded Maron and Tomás Lozano-Pérez, 'A framework for multiple-instance learning', *NeurIPS*, (1997).
- [31] Andreas Müller and Sven Behnke, 'Multi-instance methods for partially supervised image segmentation', in *Partially Supervised Learning: First IAPR TC3 Workshop, PSL 2011, Ulm, Germany, September 15-16, 2011, Revised Selected Papers 1*, pp. 110–119. Springer, (2012).
- [32] Andriy Myronenko, Ziyue Xu, et al., 'Accounting for dependencies in deep learning based multiple instance learning for whole slide imaging', in *MICCAI*, (2021).
- [33] George M Phillips, 'A de casteljau algorithm for generalized bernstein polynomials', *BIT Numerical Mathematics*, **37**, 232–236, (1997).
- [34] Linhao Qu, Xiaoyuan Luo, Shaolei Liu, Manning Wang, and Zhijian Song, 'Dgmil: Distribution guided multiple instance learning for whole slide image classification', in *MICCAI*, (2022).
- [35] Murray Rosenblatt, 'Remarks on some nonparametric estimates of a density function', *The annals of mathematical statistics*, 832–837, (1956).
- [36] Dawid Rymarczyk, Adriana Borowa, Jacek Tabor, and Bartosz Zielinski, 'Kernel self-attention for weakly-supervised image classification using deep multiple instance learning', in *WACV*, (2021).
- [37] Dawid Rymarczyk, Adam Pardy, et al., 'Protomil: Multiple instance learning with prototypical parts for whole-slide image classification', in *ECML PKDD*, (2023).
- [38] Ramprasaath R Selvaraju, Michael Cogswell, et al., 'Grad-cam: Visual explanations from deep networks via gradient-based localization', in *ICCV*, (2017).
- [39] Xiaoshuang Shi, Fuyong Xing, Yuanpu Xie, Zizhao Zhang, Lei Cui, and Lin Yang, 'Loss-based attention for deep multiple instance learning', in *AAAI*, (2020).
- [40] Korsuk Sirinukunwattana, Shan E Ahmed Raza, et al., 'Locality sensitive deep learning for detection and classification of nuclei in routine colon cancer histology images', *IEEE transactions on medical imaging*, **35**(5), 1196–1206, (2016).
- [41] Li Sun, Yanping Lu, Kaitao Yang, and Shaozi Li, 'Ecg analysis using multiple instance learning for myocardial infarction detection', *IEEE transactions on biomedical engineering*, **59**(12), 3348–3356, (2012).
- [42] Aad W Van der Vaart, *Asymptotic statistics*, volume 3, Cambridge university press, 2000.
- [43] Sudheendra Vijayanarasimhan and Kristen Grauman, 'Keywords to visual categories: Multiple-instance learning for weakly supervised object categorization', in *CVPR*, (2008).
- [44] Xinggang Wang, Yongluan Yan, Peng Tang, Xiang Bai, and Wenyu Liu, 'Revisiting multiple instance neural networks', *Pattern Recognition*, **74**, 15–24, (2018).
- [45] Yongluan Yan, Xinggang Wang, Xiaojie Guo, Jiemin Fang, Wenyu Liu, and Junzhou Huang, 'Deep multi-instance learning with dynamic pooling', in *ACML*, (2018).
- [46] 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).
- [47] Jin-Gang Yu, Zihao Wu, et al., 'Prototypical multiple instance learning for predicting lymph node metastasis of breast cancer from whole-slide pathological images', *Medical Image Analysis*, 102748, (2023).
- [48] Hongrun Zhang, Yanda Meng, et al., 'Dtfd-mil: Double-tier feature distillation multiple instance learning for histopathology whole slide image classification', in *CVPR*, (2022).
- [49] Zhendong Zhao, Gang Fu, et al., 'Drug activity prediction using multiple-instance learning via joint instance and feature selection', in *BMC bioinformatics*, volume 14, pp. 1–12. BioMed Central, (2013).
- [50] Zhi-Hua Zhou, 'Multi-instance learning: A survey', *Department of Computer Science & Technology, Nanjing University, Tech. Rep.*, **1**, (2004).
- [51] Zhi-Hua Zhou, Kai Jiang, and Ming Li, 'Multi-instance learning based web mining', *Applied intelligence*, **22**, 135–147, (2005).
- [52] Xinliang Zhu, Jiawen Yao, Feiyun Zhu, and Junzhou Huang, 'Wsa: Making survival prediction from whole slide histopathological images', in *CVPR*, (2017).
- [53] Ryszard Zielinski, *Optimal Quantile Estimators Small Sample Approach*, Polish Academy of Sciences. Institute of Mathematics, 2004.