



Pixel-Level Explanation of Multiple Instance Learning Models in Biomedical Single Cell Images

Ario Sadafi^{1,2}, Oleksandra Adonkina^{1,3}, Ashkan Khakzar², Peter Lienemann¹,
Rudolf Matthias Hehr¹, Daniel Rueckert⁴, Nassir Navab^{2,5},
and Carsten Marr¹(✉)

¹ Institute of AI for Health, Helmholtz Zentrum München - German Research Center for Environmental Health, Neuherberg, Germany

carsten.marr@helmholtz-munich.de

² Computer Aided Medical Procedures (CAMP), Technical University of Munich, Munich, Germany

³ Faculty of Mathematics, Technical University Munich, Munich, Germany

⁴ Artificial Intelligence in Healthcare and Medicine, Technical University of Munich, Munich, Germany

⁵ Computer Aided Medical Procedures, Johns Hopkins University, Baltimore, USA

Abstract. Explainability is a key requirement for computer-aided diagnosis systems in clinical decision-making. Multiple instance learning with attention pooling provides instance-level explainability, however for many clinical applications a deeper, pixel-level explanation is desirable, but missing so far. In this work, we investigate the use of four attribution methods to explain a multiple instance learning models: GradCAM, Layer-Wise Relevance Propagation (LRP), Information Bottleneck Attribution (IBA), and InputIBA. With this collection of methods, we can derive pixel-level explanations on for the task of diagnosing blood cancer from patients' blood smears. We study two datasets of acute myeloid leukemia with over 100 000 single cell images and observe how each attribution method performs on the multiple instance learning architecture focusing on different properties of the white blood single cells. Additionally, we compare attribution maps with the annotations of a medical expert to see how the model's decision-making differs from the human standard. Our study addresses the challenge of implementing pixel-level explainability in multiple instance learning models and provides insights for clinicians to better understand and trust decisions from computer-aided diagnosis systems.

Keywords: Multiple instance learning · Pixel-level explainability · Blood cancer cytology

A. Sadafi and O. Adonkina—Equal contribution.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2023

A. Frangi et al. (Eds.): IPMI 2023, LNCS 13939, pp. 170–182, 2023.

https://doi.org/10.1007/978-3-031-34048-2_14

1 Introduction

Healthcare systems are challenged by an increasing number of diagnostic requests and a shortage of medical experts. AI can alleviate this problem by providing powerful decision support systems that free medical experts from repetitive, tiring tasks [17]. However, explainability on all levels is required to ensure the proper working of deep learning ‘black box’ models, and to build trust for the widespread application of health AI.

For decision making that relies on the analysis of hundreds of single instances (e.g. histological patches [4] or single cells [18]), attention-based multiple instance learning (MIL) provides explainability on the instance level [5]. This allows algorithms to highlight suspicious structures in cancer tissue and retrieve prototypical, diagnostic cells in blood or bone marrow smears. In particular in cases where morphological features are unknown, it is of the highest importance to be able to inspect not only high attention instances, but also high attention pixels therein.

A number of different approaches for pixel-level explainability have been proposed and evaluated in the past. Backpropagation based methods such as layer-wise relevance propagation (LRP) [15] and guided backpropagation [24] leverage the gradient as attribution. Other methods work with latent features, including GradCAM [21], which utilizes the activations on the final convolution layers, or IBA [20], which measures the predictive information of latent features. These methods are widely used in the medical field to provide some level of explainability: Böhle et al. [3] use LRP to explain the decisions of the neural network on brain MRIs of Alzheimer disease patients; Arnaout et al. [2] propose an ensemble neural network to detect prenatal complex congenital heart disease and use GradCAM to explain the decisions of their expert-level model. Another attribution method, InputIBA [27], has proven to be useful for generating saliency maps for dermatology lesions [11].

Unfortunately, most of these approaches cannot be applied to MIL out of the box. Complex gradient flows and the additional dimension introduced by the bag structure in the MIL model architecture requires adapting explainability algorithms accordingly. Here, we introduce MILPLE, the first multiple instance learning algorithm with pixel-level explainability. We showcase MILPLE (Fig. 1) on two clinical single cell datasets with high relevance for the automatic classification of leukaemia subtypes from patient samples. We adapt GradCAM, LRP, IBA, and InputIBA to a MIL architecture and study the effectiveness of these methods in providing pixel-level explainability for instances. Although the quality of some of the methods seems visually plausible, quantitative analysis shows that there is no silver bullet addressing all challenges. With widespread applications of attention based MIL in different medical tasks, MILPLE helps provide pixel-level explanation using the mentioned algorithms. To foster reproducible research, our code is available on Github <https://github.com/marrlab/MILPLE>.

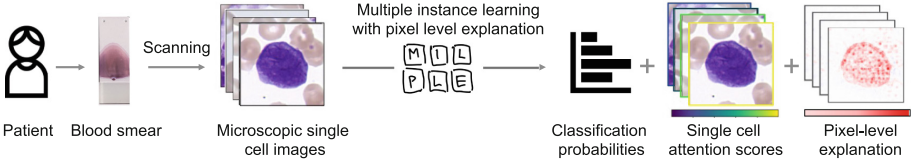


Fig. 1. MILPLE brings pixel-level explainability to multiple instance learning models. We apply MILPLE to two clinical single-cell datasets and showcase its explanatory power for revealing morpho-genetic correlations in blood cancer. In our example, blood smears from over 300 patients suffering from an aggressive leukemic subtype called acute myeloid leukemia (AML) have been digitized and microscopic images of white blood cells have been extracted. AML subtypes are predicted based on the pool of cells, and most important cells are identified based on the MIL attention mechanism, while the most important pixels in each of those are indicated with MILPLE.

2 Methodology

2.1 Multiple Instance Learning

The objective of a multiple instance learning (MIL) model f is analyzing a bag of input instances $B = \{I_1, \dots, I_N\}$ and classifying it into one of the classes $c_i \in C$ [12]. In attention-based MIL [7], an attention score $\alpha_k \in A$, $k \in \{1, \dots, N\}$ for every instance quantifies the importance of that instance for bag classification:

$$c_i, \alpha_k = f(B). \quad (1)$$

There are two approaches to implement MIL: Instance level and embedding level MIL [25]. We focus on the embedding level MIL, where every input instance is mapped into a low dimensional space via $h_k = f_{\text{emb}}(I_k, \sigma)$ with σ being learned model parameter. By pooling information distributed between the instances, one bag is aggregated into a representative bag feature vector and used for the final classification. Attention pooling [7] provides bag level of explainability and best accuracy in many problems. MIL training can be formulated as

$$\mathcal{L}_{\text{MIL}}(\theta, \sigma) = \text{CE}(c, \hat{c}) \quad (2)$$

with $\hat{c} = f_{\text{MIL}}(H, A; \theta)$, where c is the ground truth label for the whole bag, $H = \{h_1, \dots, h_N\}$ are the embedding feature vectors of all instances and CE is the cross entropy loss. θ and σ represent learnable model parameters. Based on the attention scores $\alpha_k \in A$, the bag embedding z is calculated as a weighted average over all of the embedding feature vectors:

$$z = \sum_{k=1}^N \alpha_k h_k, \quad \text{where} \quad \alpha_k = \frac{\exp\{w^T \tanh(Vh_k^T)\}}{\sum_{j=1}^N \exp\{w^T \tanh(Vh_j^T)\}}. \quad (3)$$

The parameters V and w are learned in a semi-supervised way during training. With only bag level annotation, instances with the most probable contribution to the classification are given a higher attention score.

2.2 GradCAM

Gradient-weighted Class Activation Mapping (GradCAM) is an explanation technique leveraging the gradient information to localize the most discriminative regions of an input image for a given model prediction. It computes the gradient of the predicted class score with respect to the feature maps of the last convolutional layer and weights each feature map by the corresponding gradient to obtain the class activation map. The class activation map highlights the regions of the input image that are most relevant for the prediction. Blue parts of the map indicate no contribution and red parts indicate high contribution.

2.3 Layer-Wise Relevance Propagation

Layer-wise relevance propagation (LRP) is an explanation technique for deep neural networks which produces pixel-level decomposition of the input by redistributing relevance in the backward pass [14]. Using local redistribution rules a relevance score R_i is assigned to the input variable according to the classifier output $f(x)$:

$$\sum_i R_i^0 = \dots = \sum_j R_j^{L-2} = \sum_k R_k^{L-1} = \dots = f(x) \quad (4)$$

This backward distribution is lossless, meaning that no relevance is lost in the process while also no additional relevance is introduced at every layer L . A relevance score for every input variable R_i shows the contribution of that variable to the final outcome, which is positive or negative, depending on whether that variable supported the outcome or went against the prediction. The basic rule [14] for LRP is defined as $R_j^{L-1} = \sum_k \frac{a_j w_{jk}}{\sum_j a_j w_{jk}} R_k^L$, where w_{jk} is the weight between the j and k layers and a_j is the activation of neuron j . Epsilon rule [14] is an improvement to the basic rule by introducing a positive small ϵ value in the denominator. The ϵ will consume some of the relevance making sparser explanations with less noise. Gamma rule [14] tries to favor positive contributions more by introducing a γ coefficient on positive weights such that the impact on positive weights is controlled with it. As it increases, the effect of positive weights becomes more pronounced. ZBox rule [15] is designed for the input pixel space which is constraint to boxes.

Application to MIL. MIL architectures are a complex combination of different layer types. Fully connected layers are more often used in earlier stages in comparison with normal convolutional neural networks. We tested different combinations of rules. Based on the results and suggestions introduced by Montavon et al. [14], we decided to apply ZBox rule on the first layer for every instance, gamma rule for the feature extractor f_{emb} and epsilon rule on the attention mechanism and final classifier.

2.4 Information Bottleneck Attribution

In contrast to LRP as a back-propagation method, Information bottleneck attribution (IBA) [20] is based on information theory. IBA works by placing a bottleneck on the network to restrict the flow of information by adding noise to the features. A bottleneck on the features F at a given layer can be represented by $Z = \lambda F + (1 - \lambda)\epsilon$ where ϵ is the noise controlled by λ , a mask with the same dimensions as F and elements with values between 0 and 1. The idea is to minimize the mutual information between the input X and Z while maximizing the information between Z and target Y :

$$\max_{\lambda} I(Y, Z) - \beta I(X, Z) \quad (5)$$

Here, β is the Lagrange multiplier controlling the amount of information that passes through the bottleneck. \mathcal{L}_I is an approximation of intractable term $I(X, Z)$:

$$I(X, Z) \approx \mathcal{L}_I = E_F[D_{KL}(P(Z|F) \parallel Q(Z))], \quad (6)$$

where $Q(Z)$ is a normal distribution with estimated mean and variance of F from a batch of samples. Intuitively, $I(Y, T)$ is equivalent to accurate predictions. Thus instead of maximizing it, we can minimize the loss function, cross entropy loss in our case, and therefore information bottleneck can be obtained by using $\mathcal{L} = \beta\mathcal{L}_I + CE$ as the objective.

2.5 Input Information Bottleneck Attribution

The motivation behind InputIBA [27] is to make the information bottleneck optimization in Eq. 5 possible on the input space. IBA as proposed in Eq. 5 and 6 results in an overestimation of mutual information as the bottleneck is applied on earlier layers. The formulation is the most valid when the bottleneck is applied to a deep layer where the Gaussian distribution approximation of activation values is valid [20]. Thus InputIBA proposes a trick where the optimal bottleneck is first computed using Eq. 5. Let us refer to it as Z^* . Then we look for an input bottleneck Z_G that induces the same optimal bottleneck on the deep layer. In order to make the input bottleneck Z_G induce Z^* in deep layers, the following distribution matching is done:

$$\min_{\lambda_G} D[P(f(Z_G)) \parallel P(Z^*)] \quad (7)$$

By optimizing Eq. 7 we find the optimal input bottleneck Z_G^* that induces Z^* in the selected deep layer. InputIBA proceeds to use Z_G^* as a prior for solving the information bottleneck optimization (Eq. 5). The input bottleneck Z_I is conditioned on Z_G as follows: $Z_I = \Lambda Z_G + (1 - \Lambda)\epsilon$, where Λ is the input mask. The final mask Z_I^* is computed by optimizing Eq. 5 on Z_I , and it restricts the flow in the deep layers within limits defined by Z_G^* .

Application to MIL. We had to overcome an obstacle of additional dimension introduced by the bag instances compared to conventional neural networks to apply InputIBA to the MIL structure. In comparison to standard neural networks working with single images, in MIL it is not straightforward to form a batch of bags as convolutions won't handle five dimensions. It is suggested to apply IBA on the deepest layer of the network, however in MIL architectures it seemed that applying IBA on earlier layers yields a better result. After conducting experiments and testing every convolution layers of the resnet backbone, we decided to place the bottleneck at the third convolutional layer where we obtained the best signal compared to other layers. The distance in Eq. 7 is minimized based on an adversarial optimization scheme [27]. The generative adversarial network is trained for each instance in the bag individually. We used $\beta = 40$ to control the amount of information passing through the input bottleneck.

2.6 Quantitative Evaluation of Pixel-Wise Explainability Methods

There is extensive literature studying the quality of the explanations [1, 6, 8, 9, 16, 23], but only few quantitative approaches exist. The intuition behind these methods is perturbation of features found to be important and measuring their impact on output to evaluate the quality of the feature attributions.

Insertion/Deletion [19]. Insertion method gradually inserts pixels into the baseline input (zeros) while deletion method removes pixels from input data by replacing them with the baseline value (zero) according to their attribution scores from high to low. While computing the output of the network over different percentage of insertion or deletion a curve is obtained. The area under the curve (AUC) is calculated for every input and averaged over the whole dataset. A higher AUC in insertion means important pixels were inserted first while a lower AUC in deletion means important pixels were removed first.

Remove-and-Retrain [6]. (ROAR) is an empirical measure to approximate the quality of feature attributions by verifying the degradation of the accuracy of a retrained model when the features identified as important are removed from the dataset. The processes is repeated with various percentages of removal. A sharper degradation of the accuracy demonstrates a better identification of important features. Random assignment of importance is defined as a baseline.

3 Experiments

3.1 Dataset

We study the effectiveness of pixel attribution methods on acute myeloid leukemia (AML) subtype recognition tasks using two different datasets: Deep-APL and an in-house AML dataset.

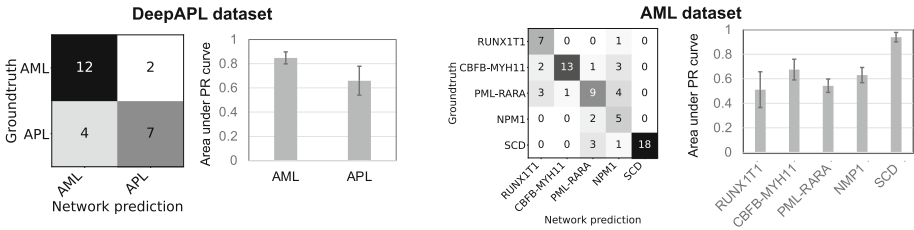


Fig. 2. The confusion matrix and area under the precision recall curve is reported for the two datasets MIL model was trained on. Mean and standard deviation are calculated over 5 independent runs.

DeepAPL [22] is a single cell blood smear dataset consisting of 72 AML and 34 acute promyelocytic leukemia (APL) patients collected at the Johns Hopkins Hospital.

AML dataset is a cohort of 242 patient blood smears from four different prevalent AML genetic subtypes [10]: i) APL with PML::RARA mutation, ii) AML with NPM1 mutation, iii) AML with RUNX1::RUNX1T1 mutation, and iv) AML with CBFB::MYH11 mutation. A fifth group of stem cell donors (SCD) comprises only healthy individuals and is thus used as the control group. Each blood smear contains at least 150 single white blood cell images resulting in a total of 81,214 cells. This dataset is available via TCIA¹.

3.2 Implementation Details

For the backbone of our approach and feature extraction from single cell images, we use the ResNeXt [26] architecture suggested by Matek et al. [13], which is pretrained on the relevant task of single white blood cell classification. Features are extracted from the last convolutional layer of the ResNeXt and passed into the MIL architecture with a second feature extraction step consisting of two convolutional layers with adaptive max-pooling and a fully connected layer. The attention mechanism consists of two fully connected layers and finally, the classifier consists of two fully connected layers. Adam Optimizer with a learning rate of 5×10^{-4} for DeepAPL and 5×10^{-5} for the AML dataset with a Nesterov momentum of 0.9 was used. The datasets are split into stratified subsets for train, validation and test in a 60-20-20 percent regime.

3.3 Model Training

The training of the MIL model on the two datasets continues for 40 and 150 epoches, respectively, while the validation loss is monitored. If the validation loss does not decrease for 5 consecutive epochs the training is stopped. We conducted 5 independent runs to train the model. Table 1 shows the mean and standard deviation of accuracy, macro F1 score and area under ROC curve.

¹ <https://doi.org/10.7937/6ppe-4020>.

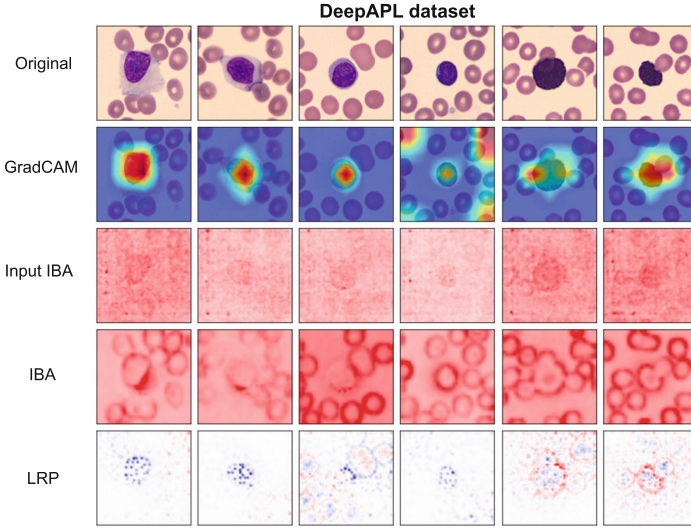


Fig. 3. Pixel-level explanation methods applied to exemplary images from the DeepAPL dataset. For GradCAM blue parts of the map indicate no contribution and red parts indicate high contribution, and similarly, for LRP blue parts indicate negative contribution and red parts indicate positive contribution. In many cases GradCAM and LRP focus on the white blood cells in the center of the image, while IBA focuses also on the red blood cell surrounding it. InputIBA shows a relatively scattered focus. (Color figure online)

3.4 Evaluation of Explanations

Qualitative Evaluation includes inspection of single cell images and comparison with medical expert annotation. Figure 3 and 4 show selected cells from both datasets and pixel-level explanations provided by the four different methods. In Fig. 4, we compare pixel attributions with expert annotations as a medical expert has annotated a small subset of single cells in the AML dataset. Most of the methods detect morphological features defined by the expert as important.

Quantitative Evaluation of the explanations is an essential step for correct understanding of what model focuses on. In order to evaluate the quality of different methods, we performed Insertion/Deletion and ROAR experiments on each of the GradCAM, LRP, IBA, and InputIBA methods as shown in Fig. 5. The performance of the method is highly dependent on the dataset and each time different methods end up to be the most suitable.

3.5 Discussion and Results

Model Performance: We compare our training on DeepAPL with the state-of-the-art method proposed by Sidhom et al. [22] on the dataset. Since the datasets

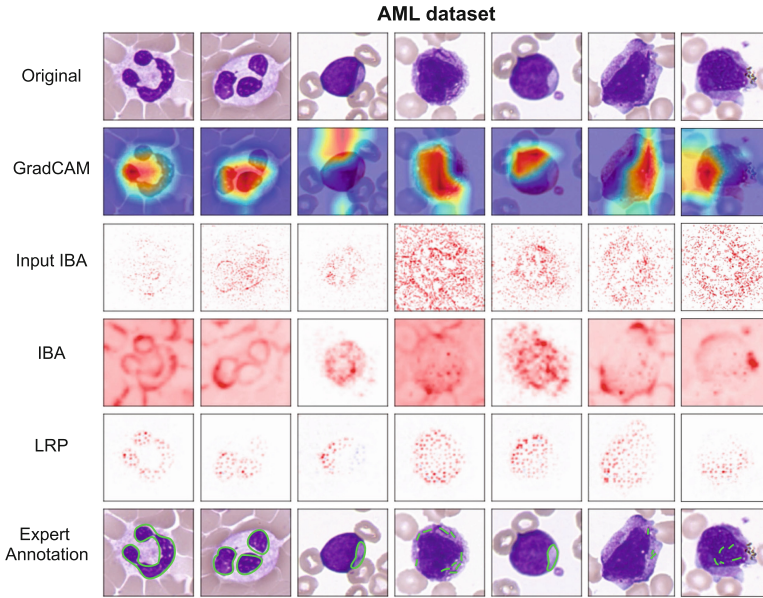


Fig. 4. Pixel-level explanation methods applied on exemplary images from AML dataset. In the first two images, all methods agree on the morphology found relevant by the expert (last row). In the following images, the methods highlight different regions and are only sometimes in concordance with the expert.

are imbalanced, we are reporting the area under precision recall curve for each class as well as the confusion matrix for both datasets to get a better view over the class-wise performance. Figure 2 shows that classification results are robust across the two datasets. On DeepAPL, with no special tailoring of the method to the dataset, we could outperform the state-of-the-art method based on sample analysis cell by cell. MIL takes all cells into consideration and can thus achieve a higher accuracy in the task. On the AML dataset some ambiguity exists between different malignant classes, which is to be expected since AML subtype classification based on cell morphology only is a challenging task even for the medical experts. Model identifies the majority of benign stem cell donors correctly.

Explanations: A close inspection of the pixel explanations from the four different methods reveals fundamental differences (see Fig. 3, 4): For the DeepAPL dataset (Fig. 3) we observe that GradCAM focuses on the white blood cell nucleus in most cases. In some cells however it fails to recognise the cell and instead puts high relevance on background pixels at the image border. Though according to our ROAR results (Fig. 5), removing the white blood cell affects the accuracy significantly, pointing to the fact that the network is using features relevant to them. InputIBA puts most focus on the centre of the image, and thus correctly on the white blood cell. However, pixel attention is spread out

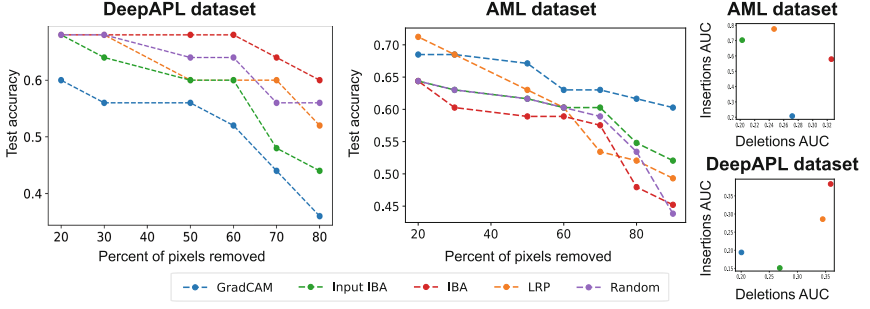


Fig. 5. Remove-and-Retrain (ROAR) experiment (left) and insertion/deletion experiment (right) for both datasets. GradCAM has the best pixel attribution in DeepAPL ROAR experiments, while on the AML dataset, LRP and IBA perform best. Insertion/deletion experiments for DeepAPL support GradCAM. For the AML dataset InputIBA and LRP have the best performance in insertion/deletion experiment.

Table 1. Mean and standard deviation of accuracy, macro F1 score and area under ROC curve is reported for the two blood cancer datasets. Our attention based MIL method outperforms the original DeepAPL method [22].

Data	Method	Accuracy	F1 score	AU ROC
DeepAPL	ours	0.65 ± 0.07	0.63 ± 0.08	0.750 ± 0.078
DeepAPL	Sidhom et al. [22]	–	–	0.739
AML	ours	0.68 ± 0.03	0.65 ± 0.04	0.855 ± 0.037

over the whole image at times (Fig. 3). The ROAR results for InputIBA (Fig. 5) also show that the accuracy drops if corresponding image regions are removed.

On the AML dataset we observe that IBA highlights image regions that correspond to either abnormal cytoplasm (4th, 6th and 7th cell from left, Fig. 4) or to structures in the nucleus (first two cells in Fig. 4). These are particularly interesting since they show that the method is able to retrieve morphological details that escape the human eye (3rd cell: the cell appears to be dark violet in the original images, but IBA is able to focus on morphology therein) and to segment granules, whose structure is relevant for cell type classification (4th cell). The ROAR results from the AML dataset (Fig. 5) show that removing morphological features identified by IBA significantly disrupts accuracy. This signifies that the model relies on these pixel during training. LRP focuses on the white blood cell in the image center and the nucleus therein. We observe that the ROAR results for LRP are not very informative (Fig. 5), and the method performs similarly to random. This might be due to the LRP structure and a problem with the ROAR metric. However, LRP achieves a good score on the Deletion/Insertion metric (Fig. 5). This means that LRP features have an immediate effect on the output of the network.

4 Conclusion

Incorporating pixel-level explainability in multiple instance learning allows us to inspect instances, evaluate the focus of our model, and find morphological details that might be missed by the human eye. All four pixel-level explainability methods we used revealed interesting insights and highlighting morphological details that fit prior expert knowledge. However, more work has to be done on systematically comparing and quantifying clinical expert annotations with explainability predictions, to eventually select appropriate methods for the application at hand, and potentially reveal novel morpho-genetic correlations.

We believe that our study will be instrumental for multiple instance learning applications in health AI. Single-cell data is ideal for method development, since it allows a direct comparison of model prediction and human intuition. However, applied to computational histopathology, where a large amount of digitized data exists, the pixel-level insight into tissue structure at multiple scales might reveal morphological properties previously unrecognized. With novel spatial single-cell RNA sequencing technologies being on the brink of becoming available widely, we expect a high demand for methods like MILPLE.

References

1. Adebayo, J., Gilmer, J., Muelly, M., Goodfellow, I., Hardt, M., Kim, B.: Sanity checks for saliency maps. In: *Advances in Neural Information Processing Systems*, vol. 31 (2018)
2. Arnaout, R., Curran, L., Zhao, Y., Levine, J.C., Chinn, E., Moon-Grady, A.J.: An ensemble of neural networks provides expert-level prenatal detection of complex congenital heart disease. *Nat. Med.* **27**(5), 882–891 (2021)
3. Böhle, M., Eitel, F., Weygandt, M., Ritter, K.: Layer-wise relevance propagation for explaining deep neural network decisions in MRI-based Alzheimer’s disease classification. *Front. Aging Neurosci.* **11**, 194 (2019)
4. Campanella, G., et al.: Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nat. Med.* **25**(8), 1301–1309 (2019)
5. Hehr, M., et al.: Explainable AI identifies diagnostic cells of genetic AML subtypes. *PLOS Digit. Health* **2**(3), e0000187 (2023)
6. Hooker, S., Erhan, D., Kindermans, P.J., Kim, B.: A benchmark for interpretability methods in deep neural networks. In: *Advances in Neural Information Processing Systems*, vol. 32 (2019)
7. Ilse, M., Tomczak, J., Welling, M.: Attention-based deep multiple instance learning. In: *International Conference on Machine Learning*, pp. 2127–2136. PMLR (2018)
8. Khakzar, A., Baselizadeh, S., Khanduja, S., Rupprecht, C., Kim, S.T., Navab, N.: Neural response interpretation through the lens of critical pathways. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13528–13538 (2021)
9. Khakzar, A., Khorsandi, P., Nobahari, R., Navab, N.: Do explanations explain? Model knows best. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10244–10253 (2022)

10. Khoury, J.D., et al.: The 5th edition of the world health organization classification of haematolymphoid tumours: myeloid and histiocytic/dendritic neoplasms. *Leukemia* **36**(7), 1703–1719 (2022)
11. Krammer, S., et al.: Deep learning-based classification of dermatological lesions given a limited amount of labelled data. *J. Eur. Acad. Dermatol. Venereol.* **36**(12), 2516–2524 (2022)
12. Maron, O., Lozano-Pérez, T.: A framework for multiple-instance learning. In: *Advances in Neural Information Processing Systems*, vol. 10 (1997)
13. Matek, C., Schwarz, S., Spiekermann, K., Marr, C.: Human-level recognition of blast cells in acute myeloid leukaemia with convolutional neural networks. *Nat. Mach. Intell.* **1**(11), 538–544 (2019)
14. Montavon, G., Binder, A., Lapuschkin, S., Samek, W., Müller, K.-R.: Layer-wise relevance propagation: an overview. In: Samek, W., Montavon, G., Vedaldi, A., Hansen, L.K., Müller, K.-R. (eds.) *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*. LNCS (LNAI), vol. 11700, pp. 193–209. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-28954-6_10
15. Montavon, G., Lapuschkin, S., Binder, A., Samek, W., Müller, K.R.: Explaining nonlinear classification decisions with deep taylor decomposition. *Pattern Recogn.* **65**, 211–222 (2017)
16. Nie, W., Zhang, Y., Patel, A.: A theoretical explanation for perplexing behaviors of backpropagation-based visualizations. In: *International Conference on Machine Learning*, pp. 3809–3818. PMLR (2018)
17. Rajpurkar, P., Chen, E., Banerjee, O., Topol, E.J.: AI in health and medicine. *Nat. Med.* **28**(1), 31–38 (2022)
18. Sadafi, A., et al.: Attention based multiple instance learning for classification of blood cell disorders. In: Martel, A.L., et al. (eds.) *MICCAI 2020*. LNCS, vol. 12265, pp. 246–256. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59722-1_24
19. Samek, W., Binder, A., Montavon, G., Lapuschkin, S., Müller, K.R.: Evaluating the visualization of what a deep neural network has learned. *IEEE Trans. Neural Netw. Learn. Syst.* **28**(11), 2660–2673 (2016)
20. Schulz, K., Sixt, L., Tombari, F., Landgraf, T.: Restricting the flow: information bottlenecks for attribution. *arXiv preprint arXiv:2001.00396* (2020)
21. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: visual explanations from deep networks via gradient-based localization. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 618–626 (2017)
22. Sidhom, J.W., et al.: Deep learning for diagnosis of acute promyelocytic leukemia via recognition of genomically imprinted morphologic features. *NPJ Precis. Oncol.* **5**(1), 1–8 (2021)
23. Sixt, L., Granz, M., Landgraf, T.: When explanations lie: why many modified BP attributions fail. In: *International Conference on Machine Learning*, pp. 9046–9057. PMLR (2020)
24. Springenberg, J.T., Dosovitskiy, A., Brox, T., Riedmiller, M.: Striving for simplicity: the all convolutional net. *arXiv preprint arXiv:1412.6806* (2014)
25. Vocaturo, E., Zumpano, E.: Dangerousness of dysplastic nevi: a multiple instance learning solution for early diagnosis. In: *2019 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pp. 2318–2323. IEEE (2019)

26. Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1492–1500 (2017)
27. Zhang, Y., et al.: Fine-grained neural network explanation by identifying input features with predictive information. arXiv preprint [arXiv: 2110.01471](https://arxiv.org/abs/2110.01471) (2021)