Attention based Multiple Instance Learning for Classification of Blood Cell Disorders

Ario Sadafi[1,2,3], Asya Makhro[4], Anna Bogdanova[4], Nassir Navab[2,5],
Tingying Peng[1,2⋆], Shadi Albarqouni[2,6⋆], and Carsten Marr[1⋆]

[1] Institute of Computational Biology, Helmholtz Zentrum Mnchen - German Research Center for Environmental Health, Germany
[2] Computer Aided Medical Procedures, Technical University of Munich, Germany
[3] Helmholtz AI, Helmholtz Center Munich, Germany
[4] Red Blood Cell Research Group, Institute of Veterinary Physiology, Vetsuisse Faculty and the Zurich Center for Integrative Human Physiology, University of Zurich, Zurich, Switzerland
[5] Computer Aided Medical Procedures, Johns Hopkins University, USA
[6] Computer Vision Lab (CVL), ETH Zurich, Switzerland

Abstract. Red blood cells are highly deformable and present in various shapes. In blood cell disorders, only a subset of all cells is morphologically altered and relevant for the diagnosis. However, manually labeling of all cells is laborious, complicated and introduces inter-expert variability. We propose an attention based multiple instance learning method to classify blood samples of patients suffering from blood cell disorders. Cells are detected using an R-CNN architecture. With the features extracted for each cell, a multiple instance learning method classifies patient samples into one out of four blood cell disorders. The attention mechanism provides a measure of the contribution of each cell to the overall classification and significantly improves the networks classification accuracy as well as its interpretability for the medical expert.

Keywords: Multiple Instance Learning · Attention · Red Blood Cells

1 Introduction

Historically, classification of hereditary hemolytic anemias, a particular class of blood disorders, is based on the abnormal shape of red blood cells. Sickle cell disease, spherocytosis, ovalocytosis, stomatocytosis: these types of anemia refer directly to the changes in cell shape, whereas genetic causes of the disease were identified later [8,7,14]. All but one (sickle cell disease) of the above-mentioned disorders are structural diseases of the red blood cell membrane caused by mutations in genes coding for the cytoskeletal proteins spectrin, ankyrin and band 3 protein, as well as protein 4.2 [7] making cells look like flowers, stars, hedgehogs, cups, droplets or spheres [2]. Characteristic changes in morphology are hallmarks


⋆ Shared senior authorship

of diseases caused by abnormalities in hemoglobin structure as for sickle cell disease and thalassemia. Somewhat more subtle are the changes in shape of red blood cells of patients harboring mutated glycolytic enzymes [9] or ion channels (like the Gardos channel or the PIEZO1 channel [18]). Independent of the cause and class of hereditary anemia, not all red blood cells but a fraction of them (often as large as 5-10% of the total cell population) are abnormally shaped. This makes diagnosis based on shape changes alone difficult, and additional tests are currently required. Furthermore, detection of abnormal shapes suffers from the subjective view of a human observer, a skillful, but possibly biased expert that may only process several hundreds of cells per patient. Instead, machine learning approaches are required and have been showing to outperform human experts in a number of clinical tasks. Classification of skin cancer at the dermatologist level proposed by Esteva et al. [6], human level recognition of blast cells by Matek et al. [15] or an AI system for breast cancer screening developed by McKinney et al. [16] are just some of the various cases machine learning excels experts. Introduction of an unbiased computer-based assessment of red blood cell shapes and their abundance in a blood sample may open new possibilities for diagnostics, assessment of disease severity, and monitoring of treatment success for the patients with rare anemias.

Multiple instance learning (MIL) is used in medical image computation when all of the instances from a patient must be taken into account and no specific label exists for each of the instances. For example, Campanella et al. [3] propose a method to whole slide pathology image classification with the MIL. Here, each whole slide image is weakly labeled as healthy or cancerous, but no specific label exists for every small image patch. Similarly, Ozdemir et al. [17] suggest a method based on MIL to classify lung cancer CT scans. In a slightly different work, Conjeti et al. [5] suggest a method of hashing for medical image retrieval based on MIL and the auxiliary branch for the vanishing gradient problem known to impede MIL approaches. While these approaches perform on a patch level, none of them is able to identify single cells that are often crucial for the diagnosis.

To this end, we propose a method based on weak patient labels and attention based MIL to classify patient blood samples into disease classes. Cells are extracted from images by an R-CNN architecture previously trained on a single cell detection task and feature maps of the backbone ResNet are passed to the proposed method. Without any cellwise labeling, the model manages to detect landmark cells for every disease in an unsupervised way by giving them the highest attention.

2   Methodology

A patients blood sample may consist of several bright-field images (see Dataset for a detailed description) and each image contains several instances / single red blood cells (see Fig. 1). A previously trained R-CNN architecture is used to find instances and extract their features (see Fig. 2). Based on these features, we propose to classify a sample into one of four diseases taking into account the fea-
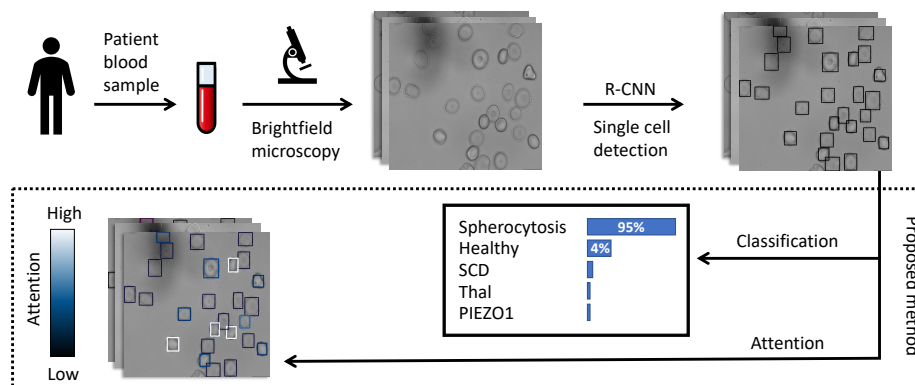
Fig. 1. Overview of the proposed method. Bright-field images of blood samples are acquired with a microscope. Using a previously trained R-CNN, all cells are detected from each image. Looking at all of the cells our proposed approach classifies the sample and provides a cell-wise attention score for better interpretability.

tures of all of the instances present in the input. Additionally, an attention score improves performance and interpretability of the method to medical experts for further verification. The proposed approach consists of three main blocks: (i) the multiple instance learning, (ii) an auxiliary single instance classifier and (iii) the attention mechanism (see Fig. 1).

More formally, our objective is a model f that classifies a blood sample containing several instances into one of the classes $c_i \in C$ and generates a score $\boldsymbol{a_k} \in A$ denoting the contribution of each instance in the final decision:

$$c_i, \boldsymbol{a_k} = f(B), \tag{1}$$

where $B = \{I_1, ..., I_N\}$ is the bag of instances and each instance I_i is a tensor of the size $256 \times 14 \times 14$.

2.1 Preprocessing: Single cell detection

Any off the shelf detection algorithm can be used for detecting single cells in the images. We are using a modified Mask R-CNN architecture [10] with ResNet[11] backbone that generates the features. For every detected cell the relevant features are extracted with RoI aligning and used as an instance I given as the input to our method. The Mask R-CNN was trained on a separate dataset from [20], consisting of 208 microscopic images with each containing 30-40 red blood cells (total of > 7000 single cells) annotated by a biomedical expert. The Mask R-CNN achieves a mean average precision (mAP) of 0.91 which is accurate enough for our application. In order to prevent cumbersome cell by cell annotation for segmentation which is not required in this approach, we limited the Mask R-CNN to a binary classification and disabled the mask head. This way, annotation of

Fig. 2. Architecture of the proposed method. The R-CNN extracts features from the input images and detects the red blood cells. Based on the detected cells a bag of cell features in all of the images is formed after RoI aligning. Passing through convolutional and fully connected layer and attention pooling a feature vector is shaped for the bag and classified. An auxiliary single instance classification (Aux-SIC) branch helps the training during the first steps. N is the number of instances in a bag and K the total number of classes.

data for single cell extraction is limited to bounding boxes around the cells that does not require any special expertise and can be performed by anyone with minimum training.

2.2   Single instance classification

Single instance classification (SIC) is the most intuitive approach for classification of samples. In our case, each instance is passed through several convolutional layers and a embedding feature vector $\boldsymbol{h}$ for every instance in the bag is generated (Fig. 2). SIC is a CNN architecture that classifies each instance embedding based on the weak labels of the bag. At inference time, a majority voting is employed to determine the class of a given sample based on the single cells' classification results:

$$\mathcal{L}_{\mathrm{SIC}}(\theta, \psi) = \frac{1}{\mathrm{N}} \sum_{\mathrm{i}=1}^{\mathrm{N}} \mathrm{CE}(\boldsymbol{c}_{\mathrm{i}}, \hat{\boldsymbol{c}}_{\mathrm{i}}), \tag{2}$$

where $\hat{\boldsymbol{c}}_{\mathrm{i}} = \mathrm{f}_{\mathrm{SIC}}(\boldsymbol{h}_{\mathrm{i}}; \psi)$, CE is the cross entropy loss, and

$$\boldsymbol{h}_{\mathrm{i}} = \mathrm{f}_{\mathrm{EMB}}(\mathrm{I}_{\mathrm{i}}, \theta) : \mathrm{I}_{\mathrm{i}} \in \mathrm{B} \tag{3}$$

with $\boldsymbol{c}_i$ being the label for each cell i based on the bag label. ψ and θ are learned model parameters.

2.3   Multiple instance learning

In contrast to supervised methods where for every given instance one tries to find a target variable $\hat{\boldsymbol{c}} \in \mathrm{C}$, multiple instance learning (MIL) tries to find the

target variable $\hat{c}$ based the input which is a set of instances $B = \{I_1, ..., I_N\}$. There are two approaches to implement MIL: instance level and embedding level approaches. We use the embedding level approach to formulate the MIL problem. As defined in eq. 3, the function f_{EMB} maps every instance into a low dimensional space $\boldsymbol{h}$ and a single representation for the whole bag z is generated using a MIL pooling method. A bag level classifier classifies z into one of the classes. A few MIL pooling methods exist. One popular method is max pooling [1] where maximization is used to generate the bag level representation:

$$z_m = \max_{k=1...N} \{h_{km}\}. \tag{4}$$

The MIL approach can be formulated as follows:

$$\mathcal{L}_{MIL}(\theta, \phi) = CE(\boldsymbol{c}, \hat{\boldsymbol{c}}), \tag{5}$$

where $\hat{\boldsymbol{c}} = f_{MIL}(\{\boldsymbol{h}_1, \ldots, \boldsymbol{h}_N\}, \{\alpha_1, \ldots, \alpha_N\}; \phi)$, where α is the attention score (see Sec. 2.4), and ϕ represents learned model parameters.

2.4 Attention mechanism

An attention mechanism is widely used in various deep learning tasks from semantic segmentation [4] to conversational question answering [23]. Ilse et al. [13] proposed an attention mechanism where a weighted average is calculated over the instance embeddings and these weights are learned by the neural network. If $H = \{h_1, ..., h_N\}$ is a bag of instance embeddings, attention based MIL pooling is defined as:

$$z = \sum_{k=1}^{N} \alpha_k h_k, \tag{6}$$

where

$$\alpha_k = \frac{\exp\{w^T \tanh(Vh_k^T)\}}{\sum_{j=1}^{N} \exp\{w^T \tanh(Vh_j^T)\}}. \tag{7}$$

V and w are parameters that are learned during the training. This attention scores α_k help the interpretability of the trained model by discovering the contribution of each instance in the drawn conclusion and acting as a similarity measure for comparison between the instances.

2.5 Overall objective function

One of the difficulties of MIL is sparse and vanishing gradients due to instance pooling. Here, we propose a dynamic loss function that incorporates the MIL loss along with the auxiliary SIC branch during the training using a decaying coefficient defined as follows:

$$\mathcal{L}(\theta, \phi, \psi) = (1 - \beta^E)\mathcal{L}_{MIL} + \beta^E \mathcal{L}_{SIC}, \tag{8}$$

where β is a hyper-parameter, and E is the epoch index.

3   Experiments

We validated the proposed method on a dataset of bright-field microscopy images of human blood cell genetic disorders. We designed an ablation study as follows: (i) single instance classification (SIC), (ii) multiple instance learning (MIL) with maxpooling, (iii) MIL with maxpooling and the auxiliary SIC branch, and (iv) MIL with attention pooling and auxiliary SIC branch.

Dataset. All images are obtained by an Axiocam mounted on Axiovert 200m Zeiss microscope with a 100x objective. No preprocessing is done and cells are not stained. The data consists of patient samples acquired at different time points, or in different kinds of solutions. In each sample there are 4 - 12 images and each image contains 12 - 45 cells. The dataset contains four genetic morphological disorders: Thalassemia (3 patients, 25 samples), sickle cell disease (9 patients, 56 samples), PIEZO1 mutation (8 patients, 44 samples) and hereditary spherocytosis (13 patients, 89 samples). Also we have a healthy control group (26 individuals, 137 samples). We did patient-wise train and test split for a fair test set selection.

Patients previously diagnosed with hereditary spherocytosis were enrolled in the CoMMiTMenT-study (`http://www.rare-anaemia.eu/`). This study was approved by the Medical Ethical Research Board of the University Medical Center Utrecht, the Netherlands, under reference code 15/426M and by the Ethical Committee of Clinical Investigations of Hospital Clinic, Spain (IDIBAPS) under reference code 2013/8436. Genetic cause of the disease was identified by the research group of Richard van Wijk, University Medical Center Utrecht, the Netherlands. The healthy subjects study was approved by the ethics committees of the University of Heidelberg, Germany, (S-066/2018) and of the University of Berne, Switzerland (2018-01766), and was performed according to the Declaration of Helsinki.

Implementation Details. The proposed method consists of three components: multiple instance embedding, auxiliary SIC, attention & bag classifier. Fig. 2 shows the architecture of the method.

The multiple instance embedding is a multi-layer convolutional neural network used for embedding features extracted by the R-CNN. It consists of five convolutional layers, a dropout layer with probability of 0.1 after the first convolution, a maxpooling layer and a fully connected layer that creates the representative 1-D embedding feature vector. These layers are common between both SIC and MIL branches and remain trainable by both branches.

An auxiliary single instance classifier looks at every instance embedding and tries to classify it with a fully connected layer. We chose a β equal to 0.5 to have a decaying contribution of this auxiliary branch during the training. Starting with a high contribution at the beginning and gradually reaching zero towards the end of the training.

Table 1. Comparison of the proposed method (MIL + att. + SIC) with other baselines. Mean and standard deviation for accuracy, weighted F1 score and average of area under ROC curve of all classes for five runs is shown.

Method	Accuracy	F1 Score	AU ROC
SIC	0.50 ± 0.01	0.46 ± 0.01	0.743 ± 0.005
MIL + max	0.46 ± 0.04	0.33 ± 0.05	0.644 ± 0.049
MIL + max + SIC	0.70 ± 0.01	0.69 ± 0.11	0.916 ± 0.005
MIL + att. + SIC	$\mathbf{0.79 \pm 0.04}$	$\mathbf{0.78 \pm 0.01}$	$\mathbf{0.960 \pm 0.003}$

In the attention and bag classifier the matrix of embedded instance representations (H) is multiplied by the attention matrix. The attention matrix is dynamically generated based on H. After the multiplication bag classifier, a fully connected layer with softmax, does the final MIL classification.

Training. We decided to use 3-fold cross validation. Three different, independent models are trained based on each fold and performances are averaged. The models are trained by AMSGrad variation of Adam optimizer [19] with a learning rate of 0.0005 and the weight decay coefficient of 10^{-5}. Training continues for a maximum of 150 epochs with an early stopping if the MIL loss drops below a specific threshold (0.005) for five consecutive epochs. The same training parameters are used across all conducted experiments. Further details about hyper parameters and implementation details can be found in our repository under `https://github.com/marrlab/attMIL`.

Evaluation metrics. Accuracy, macro F1 score and average area under the ROC and precision recall curves are used for comparison between different approaches.

Baselines. SIC is the baseline for our approach. We compare the results with a MIL without the designed auxiliary SIC branch and a MIL with a maxpooling method to our approach.

3.1   Ablation study

All of the experiments were run for five times and the average metric with standard deviation is reported. For each of the experiments we report the accuracy, weighted F1 score and area under the ROC curve in Table 1. Additionally, in Fig. 3, the area under precision recall curve for all five classes is reported.

3.2   Qualitative assessment and interpretability

The attention mechanism allows us to have a qualitative evaluation of the model. Showing the cells contributing most to the classification can be beneficial for clinical adaptability of the model as it provides the experts with some explanation of
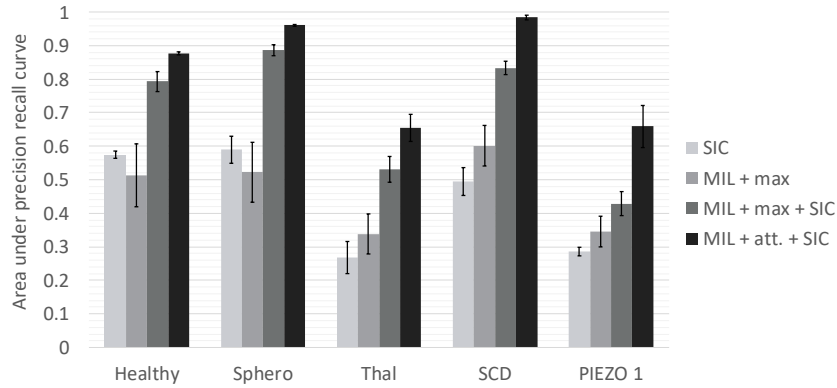
Fig. 3. Area under precision recall curve for all experiments and every class is demonstrated. We show mean and standard deviation of five runs.

Fig. 4. Exemplary whole slide images with attention values, demonstrated with colored bounding boxes. White has the highest attention score while blue and dark blue are the lowest.

the decisions made by the neural network. If cells receiving a high attention are known to be important for a specific morphological disorder, not only the model has learned them in an unsupervised manner but this also proves that the model actually knows what is relevant and what is not. Figure 4 shows cell attention in an exemplary image from samples belonging to each of the five classes.

Further, we extracted the top eight cells of a sample from every class in the dataset having the highest attention (Fig. 5). Cells are clearly morphologically different. It is interesting to note that for the healthy control class, cells that look a little bit odd received highest attentions as they might be flags for some disorders, although in the end they are not different enough to make the whole sample considered as a disorder. Note that the attention of healthy cells is generally lower than cells from disease samples (Fig. 4).

4   Conclusion

Our proposed approach based on MIL improves the performance of classification of the genetic blood disorders. The attention mechanism is effective both in

Fig. 5. Eight exemplary single cells per class with highest attention.

terms of accuracy and interpretability of classification. The model automatically learns about diagnostic cells in the samples giving them a high attention. These results are promising and have great potential for decision support and clinical applications in terms of diagnosing blood diseases and training of medical students.

Possible future works can include uncertainty estimation of the classification [22], including an active learning framework [21] and using extra features in the attention. Additionally, detailed analysis on the shape of the cells receiving high attention might be informative about the underlying pathological mechanisms and severity of the disease manifestation. This might allow stratifying patients and targeted treatments in terms of personalized medicine, which is especially important for rare anemias such as spherocytosis and xerocytosis, but also for hemolytic anemia with poor response to conventional treatment [12].

Acknowledgments

References

1. Amores, J.: Multiple instance classification: Review, taxonomy and comparative study. Artificial intelligence **201**, 81–105 (2013)
2. Bessis, M.: Corpuscles: Atlas of Red Blood Cell. Springer (1974)

3. Campanella, G., Hanna, M.G., Geneslaw, L., Miraflor, A., Silva, V.W.K., Busam, K.J., Brogi, E., Reuter, V.E., Klimstra, D.S., Fuchs, T.J.: Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. Nature medicine **25**(8), 1301–1309 (2019)

4. Chen, L.C., Yang, Y., Wang, J., Xu, W., Yuille, A.L.: Attention to scale: Scale-aware semantic image segmentation. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016)

5. Conjeti, S., Paschali, M., Katouzian, A., Navab, N.: Deep multiple instance hashing for scalable medical image retrieval. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 550–558. Springer (2017)

6. Esteva, A., Kuprel, B., Novoa, R.A., Ko, J., Swetter, S.M., Blau, H.M., Thrun, S.: Dermatologist-level classification of skin cancer with deep neural networks. Nature **542**(7639), 115–118 (2017)

7. Gallagher, P.G.: Red cell membrane disorders. ASH Education Program Book **2005**(1), 13–18 (2005)

8. Gallagher, P.: Update on the clinical spectrum and genetics of red blood cell membrane disorders. Current hematology reports **3**(2), 85–91 (2004)

9. Grace, R.F., Glader, B.: Red blood cell enzyme disorders. Pediatric Clinics **65**(3), 579–595 (2018)

10. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: Proceedings of the IEEE international conference on computer vision. pp. 2961–2969 (2017)

11. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)

12. Huisjes, R., Makhro, A., Llaudet-Planas, E., Hertz, L., Petkova-Kirova, P., Verhagen, L.P., Pignatelli, S., Rab, M.A., Schiffelers, R.M., Seiler, E., et al.: Density, heterogeneity and deformability of red cells as markers of clinical severity in hereditary spherocytosis. haematologica **105**(2), 338–347 (2020)

13. Ilse, M., Tomczak, J.M., Welling, M.: Attention-based deep multiple instance learning. arXiv preprint arXiv:1802.04712 (2018)

14. Kato, G.J., Piel, F.B., Reid, C.D., Gaston, M.H., Ohene-Frempong, K., Krishnamurti, L., Smith, W.R., Panepinto, J.A., Weatherall, D.J., Costa, F.F., et al.: Sickle cell disease. Nature Reviews Disease Primers **4**(1), 1–22 (2018)

15. Matek, C., Schwarz, S., Spiekermann, K., Marr, C.: Human-level recognition of blast cells in acute myeloid leukaemia with convolutional neural networks. Nature Machine Intelligence **1**(11), 538–544 (2019)

16. McKinney, S.M., Sieniek, M., Godbole, V., Godwin, J., Antropova, N., Ashrafian, H., Back, T., Chesus, M., Corrado, G.C., Darzi, A., et al.: International evaluation of an ai system for breast cancer screening. Nature **577**(7788), 89–94 (2020)

17. Ozdemir, O., Russell, R.L., Berlin, A.A.: A 3d probabilistic deep learning system for detection and diagnosis of lung cancer using low-dose ct scans. arXiv preprint arXiv:1902.03233 (2019)

18. Picard, V., Guitton, C., Thuret, I., Rose, C., Bendelac, L., Ghazal, K., Aguilar-Martinez, P., Badens, C., Barro, C., Bénéteau, C., et al.: Clinical and biological features in piezo1-hereditary xerocytosis and gardos channelopathy: a retrospective series of 126 patients. haematologica **104**(8), 1554–1564 (2019)

19. Reddi, S.J., Kale, S., Kumar, S.: On the convergence of adam and beyond. In: International Conference on Learning Representations (2018)

20. Sadafi, A., Koehler, N., Makhro, A., Bogdanova, A., Navab, N., Marr, C., Peng, T.: Multiclass deep active learning for detecting red blood cell subtypes in bright-

field microscopy. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 685–693. Springer (2019)

21. Settles, B., Craven, M., Ray, S.: Multiple-instance active learning. In: Advances in neural information processing systems. pp. 1289–1296 (2008)

22. Wang, X., Tang, F., Chen, H., Luo, L., Tang, Z., Ran, A.R., Cheung, C.Y., Heng, P.A.: Ud-mil: Uncertainty-driven deep multiple instance learning for oct image classification. IEEE Journal of Biomedical and Health Informatics (2020)

23. Zhu, C., Zeng, M., Huang, X.: Sdnet: Contextualized attention-based deep network for conversational question answering. arXiv preprint arXiv:1812.03593 (2018)