ALL-IDB PATCHES: WHOLE SLIDE IMAGING FOR ACUTE LYMPHOBLASTIC LEUKEMIA DETECTION USING DEEP LEARNING

Angelo Genovese, Vincenzo Piuri, Fabio Scotti

Department of Computer Science, Università degli Studi di Milano, Italy

ABSTRACT

The detection of Acute Lymphoblastic (or Lymphocytic) Leukemia (ALL) is being increasingly performed using Deep Learning models (DL) that analyze each blood sample to detect the presence of lymphoblasts, possible indicators of the disease. However, images included in current databases are either too large or already segmented. In this paper, we introduce ALL-IDB_Patches, a novel approach for processing Whole Slide Images (WSI) of ALL to take advantage of all the information available for ALL detection, by generating a larger number of samples and making the images usable by current DL models, without any pre-performed segmentation. To evaluate the attainable classification accuracy, we consider the OrthoALLNet, a Convolutional Neural Network (CNN) obtained by imposing an additional orthogonality constraint on the learned filters. The experimental results confirm the validity of our approach.

Index Terms— Deep Learning, CNN, ALL, XAI

1. INTRODUCTION

Acute Lymphoblastic (or Lymphocytic) Leukemia (ALL) refers to a disease that affects the blood cells and can, if not detected in a timely manner, spread through the body and result in fatal consequences. An important step in detecting the disease consists in inspecting the White Blood Cells (WBC) in peripheral blood samples, looking for WBCs with an altered morphology, namely lymphoblasts (Fig. 1). In fact, an excessive number of lymphoblasts in peripheral blood is one of the indicators of ALL [1, 2].

Traditionally, an experienced pathologist manually performs the analysis of WBCs looking for lymphoblasts. However, such analysis is a time consuming process that can lead to fatigue and may affect the precision of the diagnosis. Hence, to help pathologists in performing the analysis, there is a growing interest in Computer Aided Diagnosis (CAD) systems that detect the presence of lymphoblasts using image processing and Deep Learning (DL) [3]. In particular, methods based on DL are being increasingly studied for pathology analysis due to their high accuracy in medical imaging [4, 5], and have been proven to detect lymphoblasts with a high accuracy [6–8], using original learning procedures [7, 9–13], ad-hoc network architectures [14–16], or a DL-based preprocessing [6, 17].

To ensure an accurate detection of lymphoblasts using DL, it is necessary to consider datasets of labeled samples to train the model (e.g., a Convolutional Neural Network – CNN). In the case of ALL, only few datasets are available in the literature: *i)* ALL-IDB1 includes 108 Whole Slide Images (WSI) with only the position of probable lymphoblast annotated; *ii)* ALL-IDB2 includes 260 labeled WBC images [1]; *iii)* C_NMC_2019 contains 10,661 labeled WBC
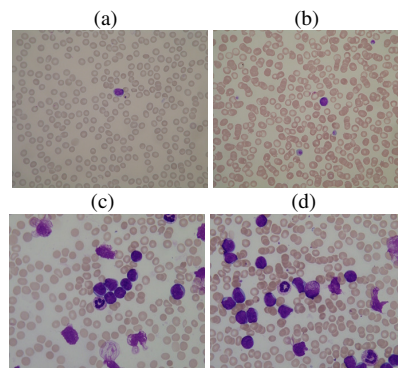
Fig. 1. Examples of images containing White Blood Cells (WBC) [1]: *(a,b) normal* WBCs; *(c,d) lymphoblasts*.

images [17]. All datasets contain images that are labeled as either *normal* or *lymphoblast*. However, these databases present some important limitations: in the case of the ALL-IDB1 dataset, the annotation does not include healthy WBCs and the images are too large to be processed by DL models. In the case of C_NMC_2019, the images have been previously segmented to only contain the area corresponding to the WBC, removing any information on the background. Moreover, both ALL-IDB2 and C_NMC_2019 databases contain images that have been previously cropped to present the WBC in the center of the image, thus influencing DL models to analyze only the central part of the image, unless data augmentations are used.

To overcome the limitations of current datasets for ALL detection, in this paper we propose the ALL-IDB_Patches[1] approach, which consists in cropping portions of the WSIs contained in the ALL-IDB1 dataset, with the purpose of making the WSIs usable for DL-based algorithms and leveraging all the information contained in WSIs. Our ALL-IDB_Patches approach has the following advantages with respect to current datasets available in the literature:

- It considers WSIs, by using the whole slide containing both red blood cells, normal WBCs, and lymphoblasts, rather than only the limited area centered around the WBC, thus taking advantage of all the information available. All the WBCs are associated with the coordinates within the WSI and the corresponding label. In particular, the WBCs are labeled in two classes: *a)* WBCs that are probable lymphoblasts; *b)* Other cases. The probable lymphoblasts have been labeled by expert oncologists in ALL-IDB1. However, some cells with a particularly altered morphology, for which it was not possible to reach a confident labeling, might have not been classified by the oncologists. The other cases have been labeled by experts in computer vision and machine learning, by focusing on stained areas not considered by the oncologists as probable lymphoblasts (Fig. 2a).

- To favor the use of DL on the proposed database, we per-

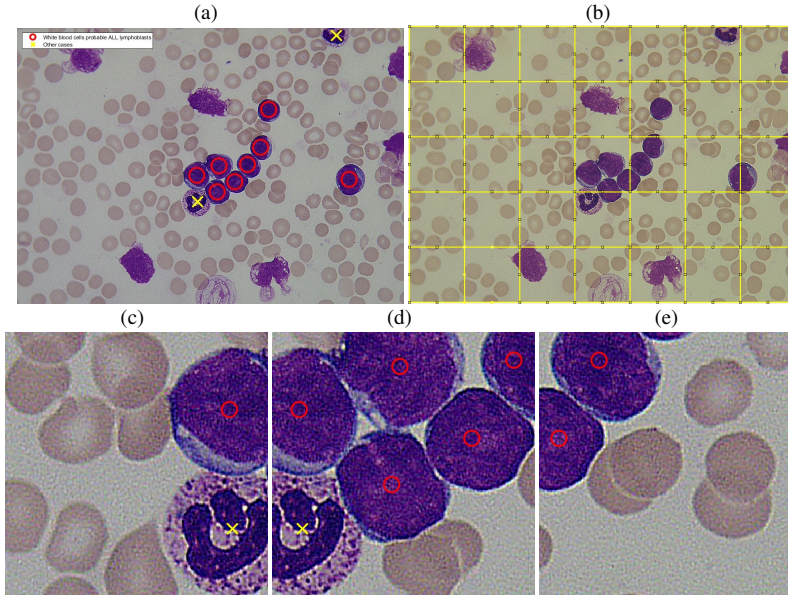[1] http://iebil.di.unimi.it/cnnALL/index.htm

Fig. 2. The ALL-IDB_Patches approach: *(a)* we consider Whole Slide Images (WSI) containing both red blood cells, normal White Blood Cells (WBCs), and lymphoblasts, rather than only the limited area centered around the WBC, thus leveraging all the information available. The position of WBCs is marked in all the WSIs, each associated with the corresponding label. The WBCs are labeled in two classes: *i)* WBCs that are probable lymphoblasts *(red circles)*; *ii)* Other cases *(yellow crosses)*. The "probable lymphoblasts" have been labeled by expert oncologists. The "other cases" have been labeled by experts in computer vision and machine learning, by focusing on stained areas not considered by the oncologists as probable lymphoblasts; *(b)* we perform a subdivision of the WSIs into patches; *(c–e)* each patch is labeled with a multi-labeling procedure, indicating whether the patch includes a WBC that is a "probable lymphoblast" *(red circles)* or an element classified as "other cases" *(yellow crosses)*.

form a subdivision of the WSIs into patches. Since resizing a WSI would cause a severe loss of information (e.g., from 2592×1944 pixels to 224×224 pixels $- \approx 1/10$ of the original size), the division of WSIs into patches enables CNN to process all the images separately. Moreover, in each patch, the WBCs can be present in different positions and sometimes not entirely contained in the patch, thus not influencing DL models to analyze only the central part of the image (Fig. 2b).

- Each patch is labeled with a multi-labeling procedure, indicating whether the patch includes a WBC that is a probable ALL lymphoblast or an element classified as "other cases". As each patch can contain both types of WBCs, the labels are not mutually exclusive. Moreover, we include the centroids for all WBCs in the patch (Fig. 2c–e).

Even though the ALL-IDB_Patches method does not produce new images with respect to ALL-IDB1, it represents a novel way of processing the WSIs contained in ALL-IDB1, with potential interest for the research communities working in ALL detection.

We also propose the first evaluation of the ALL detection accuracy on the ALL-IDB_Patches images, by describing OrthoALLNet, a novel DL methodology based on pretraining an orthogonal CNN using a multi-task learning procedure, then fine-tuning it on the ALL-IDB_Patches dataset. It is the first approach in the literature considering orthogonal CNNs for ALL detection.[2]

The remainder of the paper is structured as follows. Section 2 describes the ALL-IDB_Patches approach. Section 3 introduces the CNN-based methodology for ALL detection. Section 4 presents the experimental results. Finally, Section 5 concludes the paper.

2. THE ALL-IDB_PATCHES APPROACH

We create the images using the ALL-IDB_Patches approach by processing the WSIs of ALL-IDB1, which consist of 108 images containing about $39,000$ blood elements, including both red blood cells and WBCs. The position of WBCs is marked in all the WSIs, each labeled in one of two classes: *a)* WBCs that are probable lymphoblasts; *b)* Other cases.

The probable lymphoblasts have been labeled by expert oncologists. However, some cells with a particularly altered morphology, for which it was not possible to reach a confident labeling, might have not been classified by the oncologists. These annotations are included in ALL-IDB1. The other cases have been labeled by a team of computer vision and machine learning experts, by focusing on stained areas not considered by the oncologists. These annotations are included along with the images obtained using ALL-IDB_Patches. The annotation of the WSIs is as follows (Fig. 3).

- *Image files ImXXX_Y.jpg (from ALL-IDB1):* The image files are named with the notation *ImXXX_Y.jpg* where *XXX* is a 3-digit integer counter and *Y* is a boolean digit equal to 0 if no WBCs that are probable lymphoblasts cells are present, and equal to 1 if at least one element classified as "other cases" is present in the image. All images labeled with $Y = 0$ are from for healthy individuals, and all images labeled with $Y = 1$ are from ALL patients.

- *Text files ImXXX_Y.xyc:* Each image file *ImXXX_Y.jpg* is associated with a text file *ImXXX_Y.xyc* reporting the coordinates of the centroids of WBCs that are "probable lymphoblasts"

---

[2]We strongly discourage the use of the images obtained with the ALL-IDB_Patches approach for diagnostic or different activities than the purpose of this initiative. The images must be considered only for image processing.
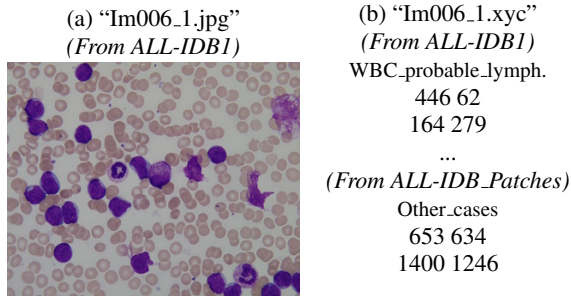
(a) "Im006_1.jpg"
(From ALL-IDB1)

(b) "Im006_1.xyc"
(From ALL-IDB1)
WBC_probable_lymph.
446 62
164 279
...
(From ALL-IDB_Patches)
Other_cases
653 634
1400 1246
... ...

Fig. 3. Example of the WSIs annotation: *(a)* image Im006_1.jpg *(from ALL-IDB1)*; *(b)* the related classification file Im006_1.xyc reporting the coordinates of the centroids of WBCs that are "probable lymphoblasts" *(from ALL-IDB1)* and the centroids of the elements classified as "other cases" *(from ALL-IDB_Patches)*.
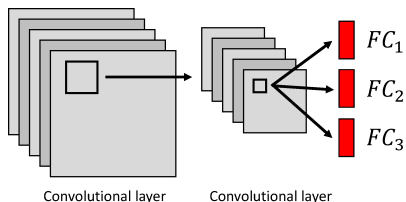

Convolutional layer   Convolutional layer

Fig. 4. The OrthoALLNet follows a multi-task learning structure, with three fully-connected layers (FC_1, FC_2, FC_3).

(from ALL-IDB1) and the centroids of the elements classified as "other cases" *(from ALL-IDB_Patches)*.

The images obtained using the ALL-IDB_Patches approach are created by cropping portions of the WSIs in ALL-IDB1, creating patches with size 256×256 pixels. To avoid any possible loss of information, we consider a 64 pixels (1/4 of the patch size) overlap between patches.[3] As a result, we obtain $10,260$ images. Each patch is associated with the corresponding label, indicating whether the patch includes a WBC that is a "probable lymphoblast" or an element classified as "other cases". As each patch can contain both types of WBCs, the labels are not mutually exclusive. Each patch is also associated with the centroids of the WBCs, if any. All the coordinates of the centroids are in the form $(x_1, y_1, x_2, y_2, ..., x_N, y_N)$, where N is the number of WBCs. Table 1 shows examples of patches and the corresponding annotations.

3. THE ALL DETECTION METHOD

This section describes the proposed method for the ALL detection using the images created using the ALL-IDB_Patches approach. Our method is based on creating the OrthoALLNet, a CNN in which the filters are learned by imposing an additional loss that maximizes the orthogonality of the filters, thus reducing feature redundancy [18]. We create the OrthoALLNet by imposing the orthogonality loss on an existing architecture (e.g., ResNet). After creating the OrthoALLNet, we perform a multi-task histopathological transfer learning [7], by pretraining the CNN on a histopathological database considering multiple labels at the same time, then fine tuning it on the images obtained using the ALL-IDB_Patches technique. The ALL detection method consists of two steps: *i)* multi-task histopathological pretraining; *ii)* ALL detection using OrthoALLNet.

3.1. Multi-Task Histopathological Pretraining

The multi-task histopathological pretraining consists of two steps. First, we create the OrthoALLNet by considering an existing

---
[3]ALL-IDB_Patches can create patches with different sizes and overlaps.

ResNet, since it represents a widely used architecture with high accuracy in several application fields [19]. We modify the last fully-connected (FC) layer to add a multi-task learning (MTL) structure consisting in three FC layers (Fig. 4).

Second, we perform the histopathological pretraining by training the OrthoALLNet on a histopathological database, composed of image patches extracted from WSIs describing several different histological tissues, such as epithelial, skeletal, adipose, or nervous. Each patch has its own label indicating which histological tissue is present in the patch, with the labels being non mutually exclusive. In particular, each patch has three different labels, organized in a hierarchical way, and each label contains a more precise indication of the kind of histological tissues present in the patch [20]. As an example, a patch p_i is associated with the following set of labels $L(p) = \{l_1, l_2, l_3\}$:

$$l_1 = \text{Epithelial (E)} ;$$
$$l_2 = \text{Simple Epithelial (E.M)} ;$$
$$l_3 = \text{Simple Squamous Epithelial (E.M.S)} . \tag{1}$$

We train the OrthoALLNet using the $L_{pretrain}$ loss, computed by aggregating the L_j losses for each FC layer and adding the L_{orth} orthogonality constraint.

1. We compute the L_j loss for the j-th FC layer by considering the corresponding l_j label. Since the labels are not mutually exclusive, we consider a multi-label soft margin loss.

2. We add the orthogonality loss L_{orth}, resulting in the following global loss $L_{pretrain}$:

$$L_{pretrain} = \left(\frac{1}{3} \sum_{j=1}^{3} L_j \right) + \lambda L_{orth} , \tag{2}$$

where $0 < \lambda < 1$ is the weight for the orthogonality loss [18].

3.2. ALL Detection using OrthoALLNet

The ALL detection consists of three steps. First, we replace the FC_1, FC_2, FC_3 layers with a single FC layer with 2 neurons, with the dimension chosen in accordance to the cardinality of the classes in ALL-IDB_Patches, in which each patch can be associated with two classes: ("probable lymphoblasts"; "Other cases").

Second, we train the OrthoALLNet by performing a deep tuning on the training subset of the images obtained using ALL-IDB_Patches, using the L_{tuning} loss. Since the classes are not mutually exclusive, we compute L_{tuning} by considering a multi-label soft margin loss L_{all}, in addition to the orthogonality loss L_{orth}:

$$L_{tuning} = L_{all} + \lambda L_{orth} . \tag{3}$$

During training, we augment the data by randomly flipping or rotating each image.

Third, we apply the trained OrthoALLNet on the testing subset of the ALL-IDB_Patches images. For each patch, the output is a binary vector indicating the absence/presence of each kind of WBCs in the image.

4. EXPERIMENTAL RESULTS

To perform the histopathological pretraining, we considered the Atlas of Digital Pathology (ADP) [20] that contains $17,668$ RGB image patches $\{p_i\}$, each with size 272×272 pixels. Each patch is associated with 3 labels, with increasing labeling precision. Each label can describe the presence of multiple histological tissues, not mutually exclusive. We train the OrthoALLNet using the procedure described in Section 3.1, with the parameters indicated in [20]. In particular, we consider $\lambda = 0.1$.

Table 1. Example of images obtained using the ALL-IDB_Patches approach and the corresponding annotations: each patch is associated with the corresponding label, indicating whether the patch includes a WBC that is a "probable lymphoblast" according to the ALL-IDB1 classification or an element classified as "other cases", and the corresponding centroids, if any. All the coordinates of the centroids are in the form $(x_1, y_1, x_2, y_2, ..., x_N, y_N)$.
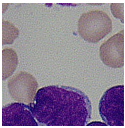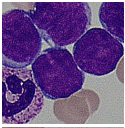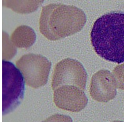
Filename	Image	Probable lymphoblast	Other cases	Centroids of "probable lymphoblasts" (x,y)	Centroids of "other cases" (x,y)
Im001_1_patch_21.jpg		1	0	(125, 226) (251, 220)	N/A
Im001_1_patch_29.jpg		1	1	(117, 149) (27, 76) (125, 34) (195, 104) (251, 28)	(30, 192)
Im001_1_patch_41.jpg		0	1	N/A	(241, 76) (15, 180)

Table 2. Accuracy results on the images obtained using the ALL-IDB_Patches approach.

Ref.	CNN	Classification Accuracy (%) (Mean$_{Std}$)
[7]	HistoTNet$_{ResNet18}$	$95.73_{0.96}$
	HistoTNet$_{ResNet34}$	$95.73_{1.08}$
[8]	ALLNet$_{ResNet18}$	$95.59_{1.00}$
	ALLNet$_{ResNet34}$	$95.46_{1.04}$
	OrthoALLNet$_{ResNet18}$	$95.91_{0.81}$
	OrthoALLNet$_{ResNet34}$	$\mathbf{96.06_{0.78}}$

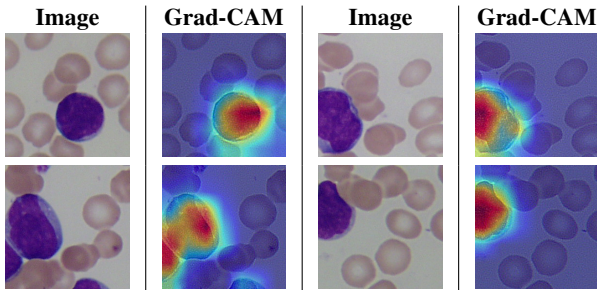| Image	Grad-CAM	Image	Grad-CAM

Fig. 5. Examples of the Grad-CAM applied on the results obtained with OrthoALLNet$_{ResNet34}$. The heatmaps are mostly overlapping with the areas of WBCs.

To perform the ALL detection, we considered a n-fold cross-validation, with $n = 5$, thus at each iteration $3/5$ of the images created using ALL-IDB_Patches are for training, $1/5$ for validation, $1/5$ for testing. After the 5 iterations, we average the results. We trained the OrthoALLNet using the Stochastic Gradient Descent (SGD) algorithm with a batch size 8, for 80 epochs, with a learning rate lr cycling in the range $[0.001, 0.02]$ every 4 epochs. Every 20 epochs, the learning rates are halved $lr' = lr/2$. After the last epoch, we use the validation subset to select the weights for which we obtained the highest classification accuracy.

Table 2 reports the classification accuracy using the proposed OrthoALLNet on the images created using the ALL-IDB_Patches

technique, using both ResNet18 and ResNet34 as architectures. As comparison, we considered the HistoTNet [7], and the ALLNet [8], since they exhibit high accuracy on the ALL-IDB2 database. All CNNs are trained using the proposed multi-task histopathological pretraining. It is possible to observe that the OrthoALLNet obtains consistently better results, in particular OrthoALLNet$_{ResNet34}$ achieves the best accuracy. The ALLNet, while exhibiting slightly inferior performance, features a reduced number of learnable parameters and could be considered in application scenarios demanding an optimized architecture, for example when performing a privacy-aware learning in edge computing [21], also in combination with neural optimization tools (e.g., OpenVino [22]). The results obtained using the methods listed in Table 2 indicate the validity of our proposed ALL-IDB_Patches approach in creating images and corresponding labels that can be used to perform an effective detection of ALL.

Fig. 5 presents the application of the Grad-CAM technique [23] on the results obtained with OrthoALLNet$_{ResNet34}$, showing how the heatmaps are mostly overlapping with the areas of WBCs, indicating the validity of our proposed methodology in learning features related to cell itself to discriminate between WBCs that are probable ALL lymphoblasts and other cases.

5. CONCLUSION

We proposed the ALL-IDB_Patches technique, which represents a novel way of processing the existing ALL-IDB1 database to make the images usable by current DL models and leverage all the information contained in Whole Slide Images (WSI) for the purpose of ALL detection. We create the images by cropping the WSIs into patches and associating each patch with labels indicating either the presence of WBCs that are "probable lymphoblasts" or "other cases". To evaluate the classification accuracy that can be obtained on the patches created using the ALL-IDB_Patches method, we consider OrthoALLNet, a CNN obtained by imposing an additional orthogonality constraint on the learned filters. The results show the validity of the proposed method in classifying the patches, establishing a baseline for future developments in the field of ALL detection. Future works will consider emerging models such as vision transformers and other histopathological transfer learning procedures.

6. REFERENCES

[1] R. Donida Labati, V. Piuri, and F. Scotti, "ALL-IDB: The Acute Lymphoblastic Leukemia Image Database for image processing," in *Proc. of ICIP*, 2011.

[2] M. M. Amin, S. Kermani, A. Talebi, and M. G. Oghli, "Recognition of acute lymphoblastic leukemia cells in microscopic images using k-means clustering and support vector machine classifier," *J. Medical Signals Sens.*, vol. 5, no. 1, Jan. 2015.

[3] H. T. Salah, I. N. Muhsen, M. E. Salama, T. Owaidah, and S. K. Hashmi, "Machine learning applications in the diagnosis of leukemia: Current trends and future directions," *Int. J. Lab. Hematol.*, vol. 41, no. 6, Dec. 2019.

[4] R. Zhang, J. Zhu, S. Yang, M. S. Hosseini, A. Genovese, L. Chen, C. Rowsell, S. Damaskinos, S. Varma, and K. N. Plataniotis, "HistoKT: Cross knowledge transfer in computational pathology," in *Proc. of ICASSP*, 2022.

[5] S. Kulkarni, N. Seneviratne, M. S. Baig, and A. H. A. Khan, "Artificial intelligence in medicine: Where are we now?," *Acad. Radiol.*, vol. 27, no. 1, Jan. 2020, Special Issue: Artificial Intelligence.

[6] A. Genovese, M. S. Hosseini, V. Piuri, K. N. Plataniotis, and F. Scotti, "Acute Lymphoblastic Leukemia detection based on adaptive unsharpening and Deep Learning," in *Proc. of ICASSP*, 2021.

[7] A. Genovese, M. S. Hosseini, V. Piuri, K. N. Plataniotis, and F. Scotti, "Histopathological transfer learning for Acute Lymphoblastic Leukemia detection," in *Proc. of CIVEMSA*, 2021.

[8] A. Genovese, "ALLNet: Acute Lymphoblastic Leukemia detection using lightweight convolutional networks," in *Proc. of CIVEMSA*, 2022.

[9] B. Masoudi, "VKCS: A pre-trained deep network with attention mechanism to diagnose acute lymphoblastic leukemia," *Multimed. Tools. Appl.*, Nov 2022.

[10] A. Loddo and L. Putzu, "On the effectiveness of leukocytes classification methods in a real application scenario," *AI*, vol. 2, no. 3, pp. 394–412, 2021.

[11] S. Shafique and S. Tehsin, "Acute Lymphoblastic Leukemia detection and classification of its subtypes using pretrained Deep Convolutional Neural Networks," *Technol. Cancer Res. T.*, vol. 17, Jan. 2018.

[12] A. Rehman, N. Abbas, T. Saba, S. I. u. Rahman, Z. Mehmood, and H. Kolivand, "Classification of Acute Lymphoblastic Leukemia using Deep Learning," *Microsc. Res. and Tech.*, vol. 81, no. 11, Nov. 2018.

[13] A. Talaat, P. Kollmannsberger, and A. Ewees, "Efficient classification of white blood cell leukemia with improved swarm optimization of deep features," *Scientific Reports*, vol. 10, no. 2536, Feb. 2020.

[14] P. Mathur, M. Piplani, R. Sawhney, A. Jindal, and R. R. Shah, "Mixup multi-attention multi-tasking model for early-stage leukemia identification," in *Proc. of ICASSP*, 2020.

[15] M. E. Billah and F. Javed, "Bayesian convolutional neural network-based models for diagnosis of blood cancer," *Applied Artificial Intelligence*, pp. 1–22, 2021.

[16] A. Kumar, J. Rawat, I. Kumar, M. Rashid, K. U. Singh, Y. D. Al-Otaibi, and U. Tariq, "Computer-aided deep learning model for identification of lymphoblast cell using microscopic leukocyte images," *Expert Systems*, p. e12894, 2021.

[17] R. Duggal, A. Gupta, R. Gupta, and P. Mallick, "SD-Layer: Stain deconvolutional layer for CNNs in medical microscopic imaging," in *Proc. of MICCAI*, 2017.

[18] J. Wang, Y. Chen, R. Chakraborty, and S. X. Yu, "Orthogonal convolutional neural networks," in *Proc. of CVPR*, 2020.

[19] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. of CVPR*, 2016.

[20] M. S. Hosseini, L. Chan, G. Tse, M. Tang, J. Deng, S. Norouzi, C. Rowsell, K. N. Plataniotis, and S. Damaskinos, "Atlas of Digital Pathology: A generalized hierarchical histological tissue type-annotated database for Deep Learning," in *Proc. of CVPR*, 2019.

[21] Y. Cai, A. Genovese, V. Piuri, F. Scotti, and M. Siegel, "IoT-based architectures for sensing and local data processing in ambient intelligence: Research and industrial trends," in *Proc. of I2MTC*, 2019.

[22] "OpenVINO," 2023, http://openvino.ai.

[23] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from Deep Networks via gradient-based localization," in *Proc. of ICCV*, 2017.