

RedDino: A foundation model for red blood cell analysis

Luca Zedda¹[0009–0001–8488–1612], Andrea Loddo¹[0000–0002–6571–3816], Cecilia Di Ruberto¹[0000–0003–4641–0307], and Carsten Marr²[0000–0003–2154–4552]

¹ University of Cagliari {luca.zedda, andrea.loddo, cecilia.dir}@unica.it

² AI for Health, Helmholtz Munich carsten.marr@helmholtz-muenchen.de

Abstract. Red blood cells (RBCs) are fundamental to human health, and precise morphological analysis is critical for diagnosing hematological disorders. Despite the potential of foundation models for medical diagnostics, comprehensive AI solutions for RBC analysis remain limited. We introduce RedDino, a self-supervised foundation model specifically designed for RBC image analysis. Leveraging a RBC-tailored version of the DINOv2 self-supervised learning framework, RedDino is trained on an extensive, meticulously curated dataset comprising over 1.25 million RBC images from diverse acquisition modalities and sources. Comprehensive evaluations demonstrate that RedDino significantly outperforms existing state-of-the-art models in the RBC shape classification. Through systematic assessments, including linear probing and nearest neighbor classification, we validate the model’s robust feature representation and strong generalization capabilities. Our key contributions are (1) a dedicated foundation model tailored for RBC analysis, (2) detailed ablation studies exploring DINOv2 configurations for RBC modeling, and (3) comprehensive generalization performance evaluation. RedDino captures nuanced morphological characteristics and represents a substantial advancement in developing reliable diagnostic tools. Source code and pre-trained models for RedDino are available at <https://github.com/Snarci/RedDino>.

Keywords: Red Blood Cell Analysis · Self-Supervised Learning · Foundation Models · Hematology · DINOv2 · Medical Imaging

1 Introduction

Hematopoiesis process is the foundation of blood, where stem cells differentiate into various blood subtypes essential for sustaining bodily functions. Understanding this intricate process is key to unraveling hematological phenomena and diseases [12]. In hematology, computer-aided diagnosis has emerged as a tool to tackle critical diagnostic challenges [25]. Two prominent areas of focus are RBC and white blood cell analysis [8, 9, 11], both of which rely extensively on imaging-based assessments to derive meaningful insights into a patient’s health. Blood smear analysis serves as the cornerstone of these investigations, involving the microscopic examination of blood smeared on glass slides. These slides

are stained to enhance cellular structures, with staining techniques tailored to specific diagnostic objectives. Variability in staining protocols and imaging acquisition introduces batch effects [32], complicating the analysis process. Beyond staining, the physical preparation of smears can also slightly alter cell morphology due to the pressure applied during the process. For efficient and robust analysis of white blood cells, the emergence of foundation models has significantly advanced the prediction of clinical outcomes [15]. These models demonstrate impressive capabilities while addressing critical challenges, such as the batch effect, a common issue in multi-source or multi-patient scenarios. In contrast, RBC analysis has yet to fully explore the potential of such technologies. This work seeks to pioneer the development of foundation models tailored for RBC analysis, addressing the nuances and requirements necessary for achieving highly representative and expressive models in this domain. Since training foundation models typically requires massive amounts of data [22], we also contributed to extracting the most complete and representative collection of RBC images in different acquisition scenarios. We demonstrate that, with carefully guided tailoring to state-of-the-art self-supervised learning methods, our family of foundation models achieves state-of-the-art performance across several RBC datasets. These models exhibit remarkable generalization capabilities, effectively mitigating the challenges posed by the batch effect. Our contributions include:

1. A dedicated family of foundation models optimized for RBC analysis, named RedDino, trained using self-supervised learning.
2. Rigorous comparative investigation studies to evaluate the effectiveness of DINOv2 [22] configurations in capturing RBC morphology.
3. Extensive benchmarking against existing state-of-the-art models, demonstrating superior performance on RBC classification and shape analysis tasks.

2 Methodology

2.1 Training Data

RedDino is trained on the largest collection of publicly available RBC image datasets to date, spanning various imaging modalities, resolutions, and staining techniques. Leveraging the DINOv2 self-supervised framework, our dataset selection was unconstrained by label availability, a significant advantage given the scarcity of annotated RBC data.

We selected 18 datasets [1–4, 6, 10, 16–20, 23, 26, 27, 29–31], comprising over 50,000 images from more than 420 individuals (Fig: 1). To mitigate the natural imbalance between red and white blood cells, we also incorporated datasets containing white blood cell images [2, 19, 26].

To extract training samples, we applied two approaches. The first involved segmentation with a fine-tuned version of CellPose [28], iteratively improved through manual corrections, producing 3,076,269 segmented cells. The second approach involved extracting non-overlapping patches of smear images with a size of 224 by 224 pixels, ensuring aspect ratio preservation and dataset diversity, generating 1,250,781 patches (Fig: 1).

Fig. 1. The RedDino training set comprises 56,712 original images. We extracted over 3 million single RBC images and more than 1.25 million patches.

Dataset	RAABIN	LDWBC	MS	Diamond	MP-IDB	T-PBS	AVI	ALL-IDB	MIMM	IML	SN-AM	NN	B	LISC	Anemia	Misc	Total
# Patients	-	72	-	10	17	4	-	-	85	-	60	106	5	-	61	-	>420
Original																	
# Images	16066	24134	1257	354	240	80	2026	368	85	345	90	10092	331	251	79	914	56712
Single cell																	
# Single cells	2176758	65238	219713	18773	24512	9471	400789	13036	6313	47683	572	38958	13503	38299	2651	26546	3076269
Patch																	
# Patches	1028224	77000	40224	19116	18480	14548	10757	7908	7480	6900	5368	4885	4739	1500	316	2500	1250781

2.2 Testing Data

To rigorously evaluate model embeddings, we used datasets with diagnostic or morphological labels as out-of-distribution test sets. Specifically, we employed the Elsafty dataset [13], the most comprehensive resource for RBC classification, containing 240,000 images across nine classes from four sources, with 72,517, 52,506, 52,103, and 63,381 images, respectively. Additional out-of-distribution test sets included the Chula dataset [21], comprising 20,875 images across 12 RBC classes, and the DSE dataset [20], which consists of 5,659 images spanning eight classes.

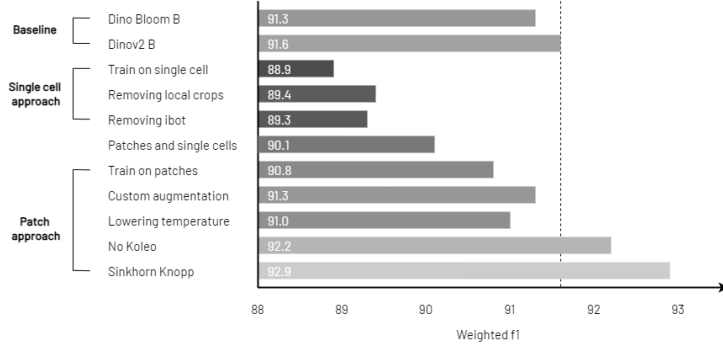
3 Experiments and Results

3.1 Model development strategies

An extensive comparative investigation was conducted to evaluate and improve RedDino’s effectiveness. Comparisons were made against baseline models, specifically DINOv2 and DinoBloom [15], the latter being the current state-of-the-art feature extractor for hematological data. The experiments were divided into two main approaches: one focusing on training with individual cells and the other on patched smear images. The optimal configuration, which guided the development of all subsequent models, was selected from the best-performing setup.

Since RedDino is designed to serve as a state-of-the-art backbone for various tasks in RBC analysis, its development was guided by evaluating feature quality in RBC shape classification, using the Elsafty dataset as the primary benchmark. Among the various training strategies explored, the DINOv2 framework was selected due to its strong performance on natural images and robust generalization across downstream tasks. The evaluation approach involved extracting features and training a linear probing to assess feature quality. Specifically, a logistic regression classifier from the sklearn library was trained on data from source 1 of the Elsafty dataset and tested on source 2. The results revealed several key findings (see Fig: 2). First, training RedDino resulted in a performance decline, similar to DinoBloom, when local crops were used in the training

Fig. 2. RedDino outperforms baseline models in linear probing evaluation by removing the Koleo regularizer and applying the Sinkhorn Knopp algorithm. The evaluation uses source 1 of the Elsafty dataset as the training set and source 2 as the test set.



process. Second, training on patched smear images instead of individual cells led to significant improvements in model performance. Third, integrating a custom augmentation pipeline, which replaced DINOv2’s pixel-level augmentations with 32-pixel-level augmentations from the Albumentations library [5], further enhanced feature quality.

Two key modifications enabled the model to outperform the state-of-the-art. The first was the removal of the Koleo regularizer, which is critical in natural image scenarios to prevent feature collapse by ensuring uniform feature distribution. However, the regularizer hindered representation quality for RBC images due to the natural uniformity of RBCs in shape and color. Pathological and abnormal RBCs, which should stand out in the feature space, were overly suppressed. The second modification involved replacing the moving average centering with the Sinkhorn Knopp centering, which improved representation quality (Figure 2).

The final model configuration, identified as the best-performing approach, is the foundation for developing the full suite of RedDino models. We proceeded to train a suite of three models for 2,000 iterations each, after which performance decreased over time, a pattern not limited to our RedDino models but a well-known phenomenon in foundation model research [24]. We finally trained our models using the same hyperparameters as the original DINOv2 while adjusting the batch size to accommodate training on two NVIDIA A100 80GB GPUs. In this setup, the RedDino small model employs a feature dimension of 384 with a batch size of 512 and contains 22 million parameters. The RedDino base model is configured with a feature dimension of 768, a batch size of 384, and 86 million parameters. Lastly, the RedDino large model features a 1,024-dimensional feature space, a batch size of 256, and comprises 304 million parameters.

3.2 Multiple scenario evaluation

Our downstream tasks aim to evaluate the expressiveness of the features extracted from the RedDino model family while keeping the evaluation scenario as

close as possible to the actual use case. We chose to focus on the classification task either because ViTs have already demonstrated impressive results in natural image segmentation and detection using adapters [7] or because state-of-the-art hematology models like CellPose [28] and RedTell [25] have achieved remarkable success.

On the classification side, the heavily imbalanced nature of the data poses a critical challenge. As with most medical imaging problems, many classes related to pathological findings are underrepresented. Therefore, we conduct experiments on the recognition of such critical RBC subtypes by training linear probing and K-nearest neighbors (K-NN) classifiers on the extracted features. The linear probing assesses the adaptability of our features to downstream tasks, while the K-NNs evaluates the effectiveness and robustness of the features under potential batch effects.

For evaluation, we employ weighted F1 measure (wF1), balanced accuracy (bAcc), and the accuracy (Acc) as metrics.

Table 1. RedDino models outperform ResNet50, DINOv2, and DinoBloom by over 2.1% in linear probing evaluation, and over 3.0% in 1-NN and 20-NN evaluation on the Elsafty dataset using a five-fold cross-validation strategy, where one source is fixed for training and the others are used for testing. "Avg Improvement" shows the average performance gain over baselines, while "Improvement" represents the performance gain compared to non-RedDino models.

Model	Linear probing			1-NN			20-NN		
	wF1	bAcc	Acc	wF1	bAcc	Acc	wF1	bAcc	Acc
ResNet50 [14]	77.6±8.1	80.3±6.0	77.8±8.0	64.3±4.8	65.8±4.7	64.2±4.9	66.2±4.9	66.9±4.9	67.5±4.5
DinoBloom-S	83.2±8.2	85.2±6.7	83.3±8.0	73.1±5.1	76.7±3.9	73.2±5.1	76.5±4.2	79.6±3.6	77.1±4.1
DinoBloom-B	84.6±6.5	85.4±6.3	84.7±6.5	72.4±6.2	75.6±5.4	72.2±6.3	76.1±6.1	78.9±5.6	76.8±5.8
DinoBloom-L	85.4±5.2	87.2±4.0	85.6±5.0	74.1±5.0	76.7±4.0	74.2±4.8	77.0±4.5	79.0±4.4	77.9±4.1
DINOv2 small	82.1±8.2	83.9±6.7	82.2±8.2	73.5±4.8	75.6±4.3	73.5±4.7	77.2±4.6	78.5±4.5	77.9±4.5
DINOv2 base	85.4±5.5	86.8±4.5	85.4±5.6	75.5±4.3	76.2±4.8	75.6±4.2	79.2±4.8	78.6±5.9	79.9±4.3
DINOv2 large	86.0±5.6	87.2±4.6	85.9±5.7	73.7±6.2	73.3±6.5	73.9±6.2	76.4±7.0	74.8±7.3	77.4±6.4
RedDino small	86.0±7.0	87.2±6.1	86.2±6.6	76.8±4.9	79.8±3.4	76.9±4.8	80.0±4.5	82.6±3.4	80.4±4.4
RedDino base	88.1±4.9	89.3±4.7	88.2±4.9	78.8±3.6	81.8±2.7	78.6±3.7	82.6±2.8	84.8±2.5	82.9±2.8
RedDino large	88.5±5.5	89.1±5.2	88.4±5.5	78.5±4.6	79.7±4.5	78.4±4.6	81.6±4.7	81.9±4.9	81.9±4.6
Improvement %	2.5	2.1	2.6	3.2	5.0	3.1	3.4	5.2	3.0
Avg improvement %	4.0	3.4	4.1	5.6	6.1	5.6	5.9	6.5	5.3

Evaluation is designed in a cross-source fashion since the Elsafty dataset is subdivided into four sources. This cross-source approach is particularly relevant in clinical trials, where significant variability is often observed in day-to-day analyses due to differences in equipment, protocols, or sample preparation methods. Our strategy involves training shallow models on the features from a fixed source and iteratively testing on the remaining three, cycling through all four possible combinations. This choice ensures a comprehensive evaluation of the model's robustness and adaptability to such variations. The results averaged across the experiments are reported in Table 1. Table 1 presents results and improvements achieved by our RedDino models compared to the selected baselines. The row *Improvement* represents the difference between the best RedDino model and the best baseline result; the row *Avg Improvement* captures

the difference between the average performance of all RedDino models and the average performance of the baselines. Often, the choice of a model for down-

Table 2. RedDino outperforms baseline models in linear probing evaluations, with the only exception of the bAcc on the DSE dataset, in a five-fold cross-validation approach. "Avg Improvement" shows the average performance gain over baselines, while "Improvement" represents the performance gain compared to non-RedDino models.


Name	Linear probing Chula			Linear probing DSE		
	wF1	bAcc	Acc	wF1	bAcc	Acc
ResNet50 [14]	76.9±0.4	70.7±0.6	76.9±0.4	83.2±0.8	51.5±9.0	83.2±0.7
DinoBloom-S	81.4±0.6	76.0±1.0	81.4±0.6	84.4±1.2	57.5±3.7	84.3±1.1
DinoBloom-B	79.0±0.4	74.3±0.5	78.9±0.4	85.3±1.0	59.9±5.4	85.4±1.0
DinoBloom-L	80.0±0.3	75.0±0.7	80.2±0.6	86.2±0.8	60.7±3.6	86.2±0.9
DINOv2 small	81.0±0.4	73.9±0.7	81.1±0.4	83.6±0.9	54.4±3.6	83.6±0.9
DINOv2 base	80.1±0.5	73.6±0.5	80.1±0.6	84.8±0.5	56.2±4.5	84.8±0.5
DINOv2 large	81.5±0.7	74.8±0.8	81.3±0.2	84.7±0.6	54.4±3.9	84.9±0.5
RedDino small	84.3±0.4	78.5±1.1	84.4±0.4	84.9±1.0	56.5±5.1	84.9±0.8
RedDino base	83.8±0.5	78.6±1.0	83.8±0.5	85.9±0.5	57.9±3.0	86.0±0.5
RedDino large	83.9±0.5	79.0±0.8	85.0±0.4	86.6±1.0	60.1±4.8	86.6±1.0
Improvement %	2.9	3.0	3.6	0.4	-0.6	0.4
Avg improvement %	4.0	3.0	4.4	1.2	1.8	1.2

stream tasks is guided by extensive and computationally expensive benchmarks designed to identify the optimal model through a brute-force search approach. In contrast, the proposed indicator offers a more realistic reflection of performance improvement. Our models demonstrated an improvement of over 2% in reference metrics for the linear probing and over 3% for the 1-NN and 20-NN, surpassing the current state-of-the-art methods.

Chula and DSE evaluation is designed to represent heavily unbalanced scenarios. We employ a 5-fold cross-validation for both datasets. For the sake of this evaluation, it is worth noting that the Chula dataset was part of the training data for the DinoBloom methods. Despite this, RedDino proves its generalization capabilities, surpassing all other models in linear probing. We do not report results for NN classification due to insufficient samples of rare classes to perform significant analysis. These strong capabilities are also evident in the DSE dataset, where RedDino outperforms all other models in almost all the selected metrics and evaluation algorithms.

It is important to note that for all analyzed datasets, the *Avg Improvement* was always positive, showcasing the importance and relevance of our models, as they consistently outperformed other representations for RBC analysis tasks across different microscopes and sources. Additionally, the evaluations yielded key insights: training RedDino on patched smear images rather than individual cells significantly enhanced model performance. RedDino base model proved to be a strong general solution, balancing performance with efficiency by utilizing fewer parameters (86 million vs. 304 million).

Fig. 3. Abnormal RBC distinguished by RedDino: the highlighted regions in (a) and (c) correlate with distinct colors in the PCA visualization (b) and (d), showcasing their differentiation provided in the embedding space. Specifically, (a) contains malaria-infected RBCs, while (c) includes echinocytes.



3.3 Visualization of the RedDino Features

PCA Visualization demonstrates feature relevance (Fig: 3). Panels (a)-(b) show a smear patch from the MP-IDB dataset with healthy and *Plasmodium falciparum*-infected RBCs and the PCA visualization from which it is observable how the model distinguishes background, cells, membranes, and parasites with high responses, highlighting the latter structures. Panels (c)-(d) apply the same approach to a brightfield image, revealing distinct features for echinocytes from the DSA dataset. This behavior arises from self-supervised training alone. The regions of interest are marked with black dotted boxes.

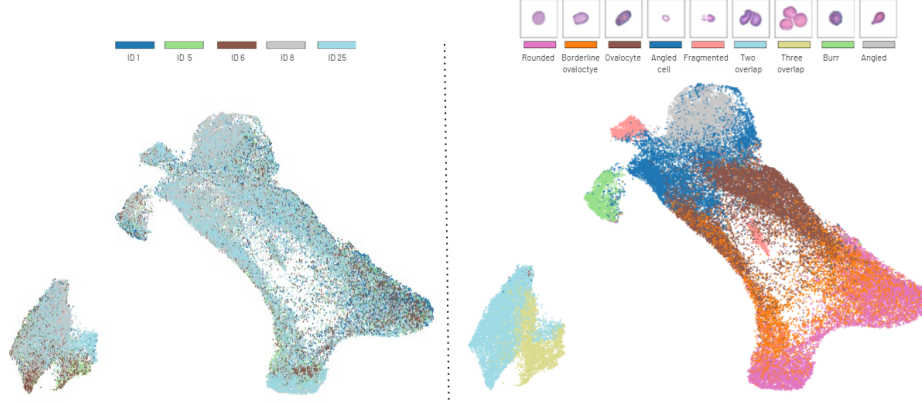
UMAP Visualization reveals batch effects within the feature space. Using the Elsafty dataset (source 1), Figure 4 shows distinct clusters without subclusters, indicating no batch across individuals. Nevertheless, overlapping classes (Rounded RBCs, Ovalocytes, and Borderline Ovalocytes) remain hard to separate due to the lack of clear clinical thresholds. Clumps form distinct clusters, confirming the model’s ability to differentiate single cells from agglomerations.

CO2 Emissions from Experiments. Our experiments were run on our in-house infrastructure, using the NVIDIA A100SXM4-80GB hardware; the total emissions are estimated to be 4.15 kg CO₂eq.

4 Conclusion

In this paper, we present RedDino, a cutting-edge foundation model specifically designed for RBC analysis. By leveraging self-supervised learning techniques, particularly a custom DINOv2 architecture, our models demonstrated superior

Fig. 4. Different classes show distinct clusters in the UMAP projection of the feature embeddings from the Elsafty dataset source 1. On the left, we show the subject distribution across the UMAP space (each patient identified by a unique identifier), while on the right, we show the class distribution.



generalization across a wide range of RBC datasets, effectively mitigating challenges such as the batch effect. The extensive range of experimentations conducted, including an extensive comparative investigation and multiple scenario evaluations, confirm that RedDino achieves state-of-the-art performance in RBC classification, surpassing previous baselines. The model’s ability to adapt to different data sources, imaging protocols, and patient populations demonstrates its robustness and potential for clinical applications. RedDino sets a new standard for RBC analysis and offers a strong foundation for future advancements in automated hematological diagnostics.

Acknowledgments. L.Z. acknowledge financial support under the National Recovery and Resilience Plan (NRRP), Mission 4 Component 2 Investment 1.5 - Call for tender No.3277 published on December 30, 2021 by the Italian Ministry of University and Research (MUR) funded by the European Union – NextGenerationEU. Project Code ECS0000038 – Project Title eINS Ecosystem of Innovation for Next Generation Sardinia – CUP F53C22000430001- Grant Assignment Decree No. 1056 adopted on June 23, 2022 by the Italian Ministry of University and Research (MUR). C.M. acknowledges funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (Grant Agreement No. 866411 & 101113551) and support from the Hightech Agenda Bayern.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Hemocompatibility of nanocrystalline diamond layers - Open Research Data - Bridge of Knowledge, <https://mostwiedzy.pl/en/open->

- research-data/hemocompatibility-of-nanocrystalline-diamond-layers,202109101719035242391-0
2. MIMM_sbilab, https://www.cancerimagingarchive.net/collection/mimm_sbilab/
 3. SN-AM, <https://www.cancerimagingarchive.net/collection/sn-am/>
 4. Arshad, Q.A., Ali, M., Hassan, S.u., Chen, C., Imran, A., Rasul, G., Sultani, W.: A Dataset and Benchmark for Malaria Life-Cycle Classification in Thin Blood Smear Images (Feb 2021)
 5. Buslaev, A., Iglovikov, V.I., Khvedchenya, E., Parinov, A., Druzhinin, M., Kalinin, A.A.: Albumentations: fast and flexible image augmentations. *Information* **11**(2), 125 (2020)
 6. Chen, H., Liu, J., Hua, C., Feng, J., Pang, B., Cao, D., Li, C.: Accurate classification of white blood cells by coupling pre-trained ResNet and DenseNet with SCAM mechanism. *BMC Bioinformatics* **23** (Jul 2022)
 7. Chen, Z., Duan, Y., Wang, W., He, J., Lu, T., Dai, J., Qiao, Y.: Vision transformer adapter for dense predictions. In: The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023 (2023)
 8. Chossegros, M., Delhommeau, F., Stockholm, D., Tannier, X.: Improving the generalizability of white blood cell classification with few-shot domain adaptation. *Journal of Pathology Informatics* p. 100405 (2024)
 9. Cuevas, E., Diaz, M., Manzanares, M., Zaldivar, D., Perez-Cisneros, M.: An Improved Computer Vision Method for White Blood Cells Detection. *Computational and Mathematical Methods in Medicine* **2013**(1), 137392 (2013)
 10. Delgado-Ortet, M., Molina, A., Alferez, S., Rodellar, J., Merino, A.: Dataset B: 331 digital images of MGG-stained blood smears from five malaria-infected patients (May 2020), <https://data.mendeley.com/datasets/2v6h4j48cx/1>
 11. Deshpande, N.M., Gite, S., Aluvalu, R.: A review of microscopic analysis of blood cells for disease detection with AI perspective <https://peerj.com/articles/cs-460>
 12. Doulatov, S., Notta, F., Laurenti, E., Dick, J.E.: Hematopoiesis: A Human Perspective. *Cell Stem Cell* **10**(2), 120–136 (Feb 2012)
 13. Elsafty, A., Soliman, A., Ahmed, Y.: 1 Million Segmented Red Blood Cells With 240 K Classified in 9 Shapes and 47 K Patches of 25 Manual Blood Smears. *Scientific Data* **11**(1), 722 (Jul 2024)
 14. He, K., Zhang, X., Ren, S., Sun, J.: Deep Residual Learning for Image Recognition. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 770–778. IEEE, Las Vegas, NV, USA (Jun 2016)
 15. Koch, V., Wagner, S.J., Kazemina, S., Sancar, E., Hehr, M., Schnabel, J.A., Peng, T., Marr, C.: Dinobloom: A foundation model for generalizable cell embeddings in hematology. In: Medical Image Computing and Computer Assisted Intervention - MICCAI 2024 - 27th International Conference, Marrakesh, Morocco, October 6-10, 2024, Proceedings, Part XII. Lecture Notes in Computer Science, vol. 15012, pp. 520–530. Springer (2024)
 16. Kouzehkanan, Z.M., Saghari, S., Tavakoli, S., Rostami, P., Abaszadeh, M., Mirzadeh, F., Satlsar, E.S., Gheidishahran, M., Gorgi, F., Mohammadi, S., Hosseini, R.: A large dataset of white blood cells containing cell locations and types, along with segmented nuclei and cytoplasm. *Scientific Reports* **12**(1), 1123 (Jan 2022)
 17. Labati, R.D., Piuri, V., Scotti, F.: All-IDB: The acute lymphoblastic leukemia image database for image processing. In: 2011 18th IEEE International Conference on Image Processing. pp. 2045–2048 (Sep 2011)

18. Loddo, A., Di Ruberto, C., Kocher, M., Prod'Hom, G.: MP-IDB: The Malaria Parasite Image Database for Image Processing and Analysis. In: Lepore, N., Brieva, J., Romero, E., Racoceanu, D., Joskowicz, L. (eds.) *Processing and Analysis of Biomedical Information*. pp. 57–65. Springer Int. Pub. (2019)
19. Madhlloom, H.T., Kareem, S.A., Ariffin, H.: Computer-aided acute leukemia blast cells segmentation in peripheral blood images. *Journal of Vibroengineering* **17**(8) (Dec 2015)
20. Makhro, A., Hegemann, I., Goede, J.S., Wijk, R.v., Mañú-Pereira, M., Sadafi, A., Marr, C., Kaestner, L., Bogdanova, A.: Red Blood Cell RedTell Dataset (Apr 2023)
21. Naruenatthanaset, K., Chalidabhongse, T.H., Palasuwan, D., Anantrasirichai, N., Palasuwan, A.: Red Blood Cell Segmentation with Overlapping Cell Separation and Classification on Imbalanced Dataset (Mar 2023)
22. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., Assran, M., Ballas, N., Galuba, W., Howes, R., Huang, P.Y., Li, S.W., Misra, I., Rabbat, M., Sharma, V., Synnaeve, G., Xu, H., Jegou, H., Mairal, J., Labatut, P., Joulin, A., Bojanowski, P.: DINOv2: Learning Robust Visual Features without Supervision (Apr 2023)
23. Rezaatfighi, S.H., Soltanian-Zadeh, H.: Automatic recognition of five types of white blood cells in peripheral blood. *Computerized Medical Imaging and Graphics* **35**(4), 333–343 (Jun 2011)
24. Roth, B., Koch, V., Wagner, S.J., Schnabel, J.A., Marr, C., Peng, T.: Low-resource finetuning of foundation models beats state-of-the-art in histopathology (Jan 2024)
25. Sadafi, A., Bordukova, M., Makhro, A., Navab, N., Bogdanova, A., Marr, C.: Redtell: an ai tool for interpretable analysis of red blood cell morphology. *Frontiers in Physiology* **14**, 1058720 (2023)
26. shenggan: Shenggan/BCCD_dataset (Jan 2025), https://github.com/Shenggan/BCCD_Dataset
27. Sidhom, J.W., Siddarthan, I.J., Lai, B.S., Luo, A., Hambley, B.C., Bynum, J., Duffield, A.S., Streiff, M.B., Moliterno, A.R., Imus, P., Gocke, C.B., Gondek, L.P., DeZern, A.E., Baras, A.S., Kickler, T., Levis, M.J., Shenderov, E.: Deep learning for diagnosis of acute promyelocytic leukemia via recognition of genomically imprinted morphologic features. *npj Precision Oncology* **5**(1), 1–8 (May 2021)
28. Stringer, C., Pachitariu, M.: Cellpose3: one-click image restoration for improved cellular segmentation (Feb 2024). <https://doi.org/10.1101/2024.02.10.579780>, <https://www.biorxiv.org/content/10.1101/2024.02.10.579780v1>
29. Sultani, W., Nawaz, W., Javed, S., Danish, M.S., Saadia, A., Ali, M.: Towards Low-Cost and Efficient Malaria Detection (Apr 2022)
30. Tyas, D.A., Ratnaningsih, T., Harjoko, A., Hartati, S.: ThalassemiaPBS (Jan 2022), <https://data.mendeley.com/datasets/gd9ysj73jd/1>
31. Vogelbacher, M., Strehmann, F., Bellaïkir, H., Mühling, M., Korfhage, N., Schneider, D., Rösner, S., Schabo, D.G., Farwig, N., Freisleben, B.: Identifying and Counting Avian Blood Cells in Whole Slide Images via Deep Learning. *Birds* **5**(1), 48–66 (Jan 2024)
32. Zedda, L., Loddo, A., Di Ruberto, C.: A deep architecture based on attention mechanisms for effective end-to-end detection of early and mature malaria parasites in a realistic scenario. *Computers in Biology and Medicine* **186**, 109704 (Mar 2025)