

Improving Sickle Cell Disease Classification: A Fusion of Conventional Classifiers, Segmented Images, and Convolutional Neural Networks

Victor Júnio Alcântara Cardoso¹, Rodrigo Moreira¹,
João Fernando Mari¹, Larissa Ferreira Rodrigues Moreira^{1,2}

¹Institute of Exacts and Technological Sciences – Federal University of Viçosa (UFV)
Rio Paranaíba – MG – Brazil

²School of Computer Science – Federal University of Uberlândia (UFU)
Uberlândia – MG – Brazil

{victor.cardoso, rodrigo, joaof.mari, larissa.f.rodrigues}@ufv.br
larissarodrigues@ufu.br

Abstract. *Sickle cell anemia, which is characterized by abnormal erythrocyte morphology, can be detected using microscopic images. Computational techniques in medicine enhance the diagnosis and treatment efficiency. However, many computational techniques, particularly those based on Convolutional Neural Networks (CNNs), require high resources and time for training, highlighting the research opportunities in methods with low computational overhead. In this paper, we propose a novel approach combining conventional classifiers, segmented images, and CNNs for the automated classification of sickle cell disease. We evaluated the impact of segmented images on classification, providing insight into deep learning integration. Our results demonstrate that using segmented images and CNN features with an SVM achieves an accuracy of 96.80%. This finding is relevant for computationally efficient scenarios, paving the way for future research and advancements in medical-image analysis.*

1. Introduction

Sickle cell anemia is a hereditary disease that results in the irregular shaping of erythrocytes, causing blockages in blood flow throughout the body. Instead of the smooth, circular shape characteristic of normal cells, individuals with sickle cell anemia have erythrocytes that are sickle-like in shape. This abnormal shape leads to a shortened lifespan of erythrocytes, resulting in anemia. Patients experience various symptoms, including pain crises due to blood vessel obstruction, fatigue, myocardial infarction, cerebrovascular accident (CVA), renal disease, pulmonary embolism, and infections. Early diagnosis, combined with appropriate interventions, significantly reduces mortality in the first five years of life from 25% to less than 3% [Hoffman et al. 2013] [Kato et al. 2018] [Kavanagh et al. 2022].

The newborn bloodspot screening test is a valuable tool for early diagnosis of sickle cell disease [Kavanagh et al. 2022]. However, the availability of this diagnostic method is limited in various emerging regions due to constraints such as insufficient financial resources, lack of laboratory supplies, and a shortage of specialized professionals.

For instance, in Brazil, 100% of hospitals in Minas Gerais conduct newborn bloodspot screening tests, while in Amapá, the coverage is only around 55% [Kato et al. 2018].

Image processing and machine learning techniques have played a crucial role in disease detection and have emerged as cost-effective solutions, particularly beneficial for emerging countries and remote regions [Schwalbe and Wahl 2020]. In this context, the automatic classification of sickle cell disease using microscopy images has gained significant attention as a promising and globally accessible alternative.

Conventional classification algorithms, also known as traditional machine learning methods, are commonly utilized in pattern recognition tasks and have found applications in various medical domains [de Faria et al. 2018, Backes 2022]. However, these algorithms rely on manual feature extraction. To overcome this limitation, Convolutional Neural Networks (CNNs) have shown remarkable results compared to traditional classifiers. The key advantage of CNNs is their ability to extract features during training, eliminating the need to preprocess the image [Goodfellow et al. 2016, Ponti et al. 2017].

In this paper, we propose to evaluate the performance of CNNs by combining them with traditional classifiers to support sickle-cell disease diagnosis using microscopy images. The main contribution of this paper relies on the direct mastery of CNN for feature extraction and feeding a classical classifier such as Support Vector Machines (SVM) and Bayes with these features, requiring lower time to train and predict while demanding low resources footprint. Additionally, our approach is suitable for retrieving relevant features from images using an Artificial Intelligence as a Service (AIaaS) Architecture [Rodrigues Moreira et al. 2023].

The remaining of this paper is organized as follows: Section 2 provides an overview of the related work conducted in the field. Section 3 details the proposed approach. The results obtained from the experiments are presented and discussed in Section 7. Finally, we conclude the paper in Section 8 and discuss potential avenues for future research.

2. Related Work

Recent advancements in Artificial Intelligence (AI) have significantly impacted medical imaging diagnostics across various domains, such as radiology, histology, oncology, and dermatology [Carmo and Mari 2021, Backes 2022, Rodrigues Moreira et al. 2023]. These advancements have improved diagnostic accuracy and efficiency, leading to the development of robust computational methods [Zhou et al. 2021]. In this section, we provide an overview of studies that have explored the utilization of AI methods for diagnosing sickle cell disease.

The classification of sickle cell disease in the *erythrocytesIDB* dataset was initially proposed by [Gual-Arnau et al. 2015]. They employed an active contour segmentation technique based on integral geometry to extract features from erythrocytes. The erythrocytes classification was performed using the K-Nearest Neighbors (KNN) algorithm.

[Rodrigues et al. 2016] presented an approach that focused on image segmentation to distinguish normal, sickle, and other deformed cells from the background in microscopy images. They applied morphological feature extraction, considering a set of nine simple geometric shape features. To identify the most discriminative features, they

employed univariate ANOVA. The study evaluated three classifiers: Naive Bayes, SVM, and KNN. The results indicated that the SVM classifier achieved the highest performance using the selected feature set, excluding the perimeter.

[de Faria et al. 2018] introduced a novel feature encoding process for red blood cell classification. They combined Scale Invariant Feature Transform (SIFT) and Speeded-Up Robust Features (SURF) key point descriptors by stacking them into a single matrix. The authors then evaluated the performance of two classification algorithms, SVM and Multi Layer Perceptrons (MLP), in classifying the red blood cells. The experimental results demonstrated that the SVM classifier achieved the best performance.

[da Silva et al. 2020] proposed a method that leveraged data augmentation techniques in conjunction with a CNN to enhance sickle cell disease classification performance. They employed the Bayesian optimization algorithm to identify an effective data augmentation strategy specifically tailored for erythrocyte images. Moreover, they utilized a lightweight CNN architecture for the classification task.

[Petrović et al. 2020] introduced an approach to effectively analyze red blood cell morphology by selecting the optimal classification method and features. Their study evaluated seven conventional classifiers, including SVM, and an optimization strategy to identify the best-performing approach. For feature extraction, the authors leveraged established methods from the literature, encompassing shape and texture features. Machine learning techniques were then employed for the task of morphology classification.

[Paz-Soto et al. 2020] presented an approach to classify sickle cell disease in microscopy images by employing neural networks trained with features extracted using integral geometry-based functions. By utilizing these functions, which capture contour information of red blood cells, the authors aimed to improve classification accuracy. However, one potential drawback of using integral geometry-based features is their limited descriptive power for capturing complex structural variations in red blood cells and sensibility to noise or small perturbations in the image, potentially leading to misclassification or reduced accuracy in some instances.

The study conducted by [Alzubaidi et al. 2020] utilized transfer learning to distinguish sickle cells from normal red blood cells across three distinct datasets: a main dataset, a transfer learning training dataset, and a testing dataset. The authors proposed three CNN architectures and achieved accuracy levels exceeding 98% by leveraging transfer learning, feature extraction, and SVM as the classifier. However, it should be noted that training a CNN with a dataset from the same domain is often infeasible in real-world scenarios.

Among the various studies reviewed, the majority have focused on the direct utilization of morphological features extracted from the contour of red blood cells. These features encompass circular and elliptical shape coefficients, roundness, eccentricity, perimeter, area, diameter, and integral geometry, among others. Notably, using shape features in conjunction with classifiers such as SVM and Neural Networks has shown promising performance.

However, these previous approaches are predominantly based on manual feature design and selection, which can introduce subjectivity and consume a significant amount of time. Additionally, this process necessitates specialized knowledge and may impede

the scalability and generalizability of the approach to different datasets. In contrast, deep learning-based classification mitigates the need for extensive preprocessing techniques. Nevertheless, the computational cost of training deep learning models can pose limitations when deploying them in real-world scenarios.

To the best of our knowledge, there is a lack of studies that utilize segmented images to classify red blood cells using the combination of conventional classifiers with deep learning techniques. Therefore, our study aims to enhance the classification of sickle cell disease by integrating traditional classifiers, segmented images, and CNNs. We also investigate the performance of our proposed approach on both original and segmented images to determine the most effective method for aiding in the diagnosis of sickle cell disease.

3. Proposed Approach

This section presents our proposed approach for improving the classification of sickle cell disease. Our main goal is to develop a comprehensive and effective method that accurately classifies sickle cell disease using a fusion of traditional and deep learning techniques. We also compared feeding the CNNs with the original and segmented images. Our study specifically focuses on identifying the best classification strategy to enhance overall performance. The steps of our proposed method are depicted in Figure 1.




Figure 1. Steps of the proposed approach.

3.1. Dataset

The images used in this study were obtained from the *erythrocytesIDB* dataset¹ provided by the University of the Balearic Islands [Gonzalez-Hidalgo et al. 2014]. The dataset comprises 626 images, each featuring a single, centrally positioned cell, categorized into three classes: normal erythrocytes (202 images), sickle cells (211 images), and erythrocytes with other deformations (213 images). All images are in JPG format and have a resolution of 80×80 pixels. Additionally, we utilized the segmented images generated using the method proposed by [Rodrigues et al. 2016]. Figure 2 shows some images from each class in the dataset, presenting both the original and segmented versions.

¹Available in: <http://erythrocytesidb.uib.es/>




Figure 2. Image instances from the *erythrocytesIDB* dataset showing original (top) and segmented (bottom) images.

4. Traditional Classifiers

The classification task is crucial in various fields, including medical imaging, where an accurate and efficient diagnosis is essential [Zhou et al. 2021, Nazir et al. 2023]. Traditional classifiers have been widely used for their interpretability and ability to handle high-dimensional data [Aljuaid and Anwar 2022]. In this study, we utilized two well-established classifiers: SVM and Naive Bayes.

The SVM is a widely used traditional classifier in machine learning due to its ability to handle complex datasets and its robustness against overfitting. It is particularly effective in handling high-dimensional data and finding optimal decision boundaries. SVM works by mapping input data into a higher-dimensional feature space and identifying a hyperplane that maximally separates different classes. It achieves this by maximizing the margin between the classes, allowing for better generalization to unseen data [Cortes and Vapnik 1995].

Naive Bayes is another popular traditional classifier, particularly known for its simplicity and efficiency [Wickramasinghe and Kalutarage 2021]. It is based on Bayes' theorem and assumes independence among the features. Despite this naive assumption, Naive Bayes often delivers competitive performance. Naive Bayes models estimate the probability of each class given the input features and then classify the data based on the highest probability. The classifier calculates the conditional probabilities using the training data and utilizes these probabilities to make predictions on unseen data. Additionally, Naive Bayes is known for its fast training and prediction times, making it suitable for various applications [Duda et al. 2000].

We selected these algorithms as our classification models due to their well-established effectiveness and their suitability for our research objective of improving sickle cell disease classification. By employing these classifiers, we aim to investigate their capabilities in accurately classifying the extracted features from the images and compare their performance against other advanced techniques in our study.

5. Evaluated Architectures

Convolutional Neural Networks (CNNs) are a powerful class of deep learning models widely used in computer vision tasks. They are specifically designed to process and analyze visual data, such as images. By leveraging hierarchical feature extraction and parameter sharing, CNNs can learn complex patterns and structures in images [Goodfellow et al. 2016][Ponti et al. 2017]. In this study, we investigated the effectiveness of three CNNs: DenseNet-169, ResNet-50, and MobileNet, chosen due to their success in previous image classification tasks.

DenseNet architecture has demonstrated excellent performance in image classification tasks [Zhou et al. 2022]. It introduces the concept of dense connections, where each layer is directly connected to every other layer in a feed-forward fashion. This connectivity pattern allows for better information flow and feature reuse, improving gradient propagation and alleviating the vanishing gradient problem [Huang et al. 2016]. We considered in this study the DenseNet-169.

ResNet, short for Residual Network, is a popular deep neural network architecture known for its residual connections. These connections enable the network to learn residual mappings, making it easier to train deeper networks. By bypassing certain layers and allowing direct information flow, ResNet can effectively tackle the challenges of training very deep neural networks, leading to improved accuracy and performance [He et al. 2016]. In this study, we evaluated the ResNet with 50 layers.

MobileNet is a lightweight deep neural network architecture designed for efficient computation and deployment on resource-constrained devices. It utilizes depth-wise separable convolutions to reduce computational complexity while maintaining good accuracy. MobileNet achieves a good balance between model size and accuracy, making it suitable for applications with limited computational resources, such as mobile devices and embedded systems [Howard et al. 2017].

We selected DenseNet-169, ResNet-50, and MobileNet for our study due to their unique structural differences as detailed in [Rodrigues Moreira et al. 2023]. Through the evaluation of these networks, our objective is to assess their specific advantages and limitations in the classification of sickle cell disease. This analysis will provide valuable insights into the impact of structural variations on both model performance and feature extraction, enhancing our understanding of the effectiveness of CNNs in this domain. We used these architectures as feature extractors by removing the classification layer at the end of training. We extracted 232,736 features from DenseNet-169, 90,947 features from ResNet-50, and 33,792 features from MobileNet.

6. Validation Protocol

All CNN models were trained and tested using the stratified k -fold cross-validation approach ($k=5$) to assess the classification performance. This involved dividing the dataset into five subsets, using one subset as the validation set while training the classifiers on the remaining subsets. By repeating this process for each subset, we obtained a comprehensive evaluation of the classifiers' performance, minimizing bias and variability in the results.

We present our findings by calculating the average accuracy, precision, recall,

and F1-score [Duda et al. 2000], which are aggregated across the 5-fold cross-validation process.

- Accuracy: measures the overall correctness of the classification (Eq. 1).

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (1)$$

- Precision: quantifies the proportion of correctly classified positive instances (Eq. 2).

$$Precision = \frac{TP}{TP + FP} \quad (2)$$

- Recall: determines the ability to identify all positive instances (Eq. 3).

$$Recall = \frac{TP}{TP + FN} \quad (3)$$

- F1-Score: provides a harmonic mean of precision and recall, offering a balanced assessment of the classifier’s performance (Eq. 4).

$$F1-Score = 2 \times \frac{Precision \times Recall}{Precision + Recall} \quad (4)$$

Where TP represents true positives, TN denotes true negatives, FP signifies false positives, and FN indicates false negatives. For each fold, we calculate these metrics and subsequently calculate the average metrics across all folds.

7. Results and Discussion

All experiments were conducted using Python (version 3.6) and implemented in Keras [Chollet et al. 2015] with TensorFlow as the backend. The classification evaluation was performed on Google Colaboratory, utilizing an Intel(R) Xeon(R) 2.30GHz processor, 12GB RAM, and the NVIDIA Tesla T4 GPU. Aiming at reproducibility, our code is available and the results in the open-source repository².

CNNs as Feature Extractors. Initially, we conducted experiments to assess the performance of CNNs as feature extractors. Initially, the original images were used without any preprocessing. The features extracted by the CNNs were then employed as feature vectors for conventional classifiers, namely Naive Bayes and linear SVM. For the linear SVM, a constant of 2.9 was empirically defined.

Table 1 presents the performance of each classifier, with the best result achieved by extracting features using the MobileNet architecture and classifying with SVM, achieving an accuracy of 91.21%. When considering feature extraction with ResNet-50, it is evident that both evaluated classifiers did not achieve satisfactory results.

Subsequently, experiments were conducted to assess the impact of segmenting images using the method proposed by [Rodrigues et al. 2016] on feature extraction. The experimental results, shown in Table 2, demonstrate that segmented images improved the

²Source code available on:

<https://github.com/larissafrodrigues/sickle-cell-classification-ENIAC2023>

Table 1. Classification performance considering CNN as feature extractor and original images.

	Accuracy (%)		Precision (%)		Recall (%)		F1-Score (%)	
	Bayes	SVM	Bayes	SVM	Bayes	SVM	Bayes	SVM
DenseNet	82.60	90.59	83.47	91.23	82.67	90.66	82.82	90.61
ResNet	59.10	68.05	61.30	70.26	59.73	67.88	52.91	64.69
MobileNet	85.30	91.21	87.48	91.42	85.30	91.28	85.39	91.07

performance of both evaluated classifiers. It is evident that when classifying segmented images, for all CNNs used as feature extractors, the Naive Bayes classifier achieved an accuracy higher than 90%, and the SVM classifier achieved an accuracy higher than 95%.

Table 2. Classification performance considering CNN as feature extractor and segmented images.

	Accuracy (%)		Precision (%)		Recall (%)		F1-Score (%)	
	Bayes	SVM	Bayes	SVM	Bayes	SVM	Bayes	SVM
DenseNet	95.37	95.84	95.74	96.00	95.38	95.87	95.44	95.88
ResNet	90.26	96.32	90.98	96.55	90.26	96.36	90.41	96.33
MobileNet	94.26	96.80	94.69	96.88	94.26	96.84	94.31	96.81

When considering classification with SVM and feature extraction using ResNet-50, the results increased from 68.05% when using original images (Table 1) to 96.32% when using segmented images (Table 2). This result suggests that ResNet-50 was able to extract important shape patterns. Finally, the best result was achieved when considering feature extraction with MobileNet and segmented images, achieving an accuracy of 96.80%.

CNN with Transfer Learning. We conducted experiments to evaluate the performance of CNNs as classifiers. We also used the trained CNNs as feature extractors by removing the top layer, and using the features to feed traditional classifiers. Once again, we evaluated the performance of each CNN when trained with original and segmented images, which were resized to 224×224 pixels (the input size allowed by each CNN). The training process was conducted by fine-tuning a CNN pre-trained with ImageNet [Krizhevsky et al. 2012] over 50 epochs, with a learning rate of 1×10^{-4} , optimization using Stochastic Gradient Descent (SGD), a momentum coefficient of 0.9, and data augmentation based on horizontal and vertical flips.

Tables 3 and 4 present the results obtained when classifying the original and segmented images, respectively. It is evident that the use of segmented images brought improvements to the ResNet-50 and MobileNet architectures. In particular, the results obtained by MobileNet demonstrate a significant improvement in all classification metrics

when using segmented images compared to original images. The accuracy increased from 87.70% with original images to 95.54% with segmented images.

Table 3. Classification performance for each CNN considering original images.

CNN	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
DenseNet	96.65	96.89	96.71	96.62
ResNet	92.65	93.13	92.70	92.56
MobileNet	87.70	88.73	87.85	87.11

Table 4. Classification performance for each CNN considering segmented images.

CNN	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
DenseNet	96.65	96.78	96.67	96.65
ResNet	95.37	93.50	95.42	95.38
MobileNet	95.54	95.62	95.59	95.54

This improvement suggests that the segmentation process proposed by [Rodrigues et al. 2016] effectively enhances the discriminative features captured by MobileNet, resulting in more accurate classifications. The deep convolutional layers of MobileNet enable it to effectively leverage the spatial information provided by the segmentation, allowing it to capture intricate patterns and structures in the images. These findings highlight the potential benefits of utilizing preprocessed images in certain scenarios, as it can significantly improve the performance of CNN-based image classification models.

On the other hand, the experimental results demonstrate that the DenseNet architecture achieved the same result for both evaluations, whether using original images or segmented images. This finding suggests that DenseNet is capable of effectively capturing important features from the original images alone, without the need for additional preprocessing, such as segmentation. This result highlights the robustness of DenseNet in handling different types of cell images, making it an alternative for sickle cell image classification tasks where preprocessing steps may not be necessary or yield significant improvements.

Figure 3 illustrates an example of feature maps extracted by DenseNet using original and segmented images. By analyzing these maps, it can be observed that DenseNet effectively captures various patterns and textures present in the original images. This demonstrates the ability of DenseNet to learn discriminative representations directly from the raw image data without the need for additional preprocessing steps.

To summarize the results, the graphs presented in Figure 4 illustrate the performance of each CNN architecture in terms of accuracy. The results are shown for two different approaches: classification with the respective CNN itself and feature extraction with SVM and Naive Bayes classifiers.


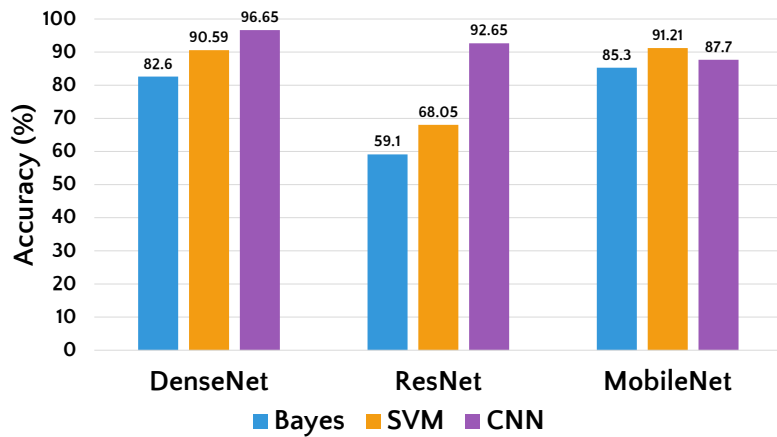
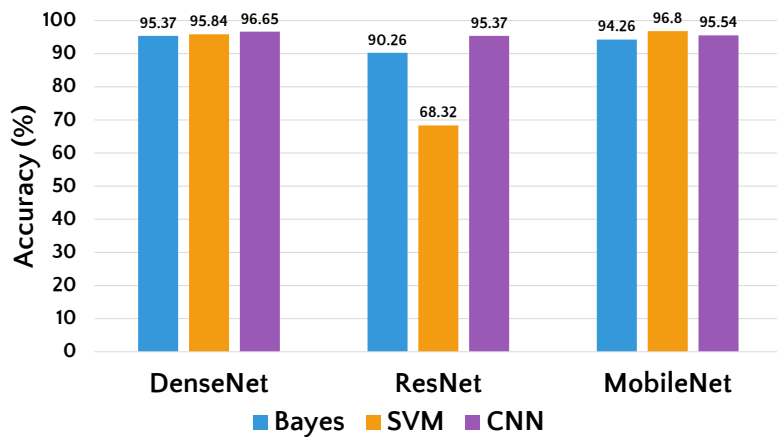


Figure 3. Examples of visualization of feature activation obtained by DenseNet.



(a) Original Images



(b) Segmented Images

Figure 4. Comparison between each strategy evaluated considering accuracy performance.

Comparison with Literature. Finally, our best result obtained for the *erythrocytesIDB* dataset was compared with other state-of-the-art approaches in the literature, as shown in Table 5. The comparison reveals that our best score is on par with or surpasses the performance of the leading techniques reported in the literature. This comparison underscores the competitiveness of our approach and highlights its effectiveness in addressing the classification of sickle cell disease in microscopy images.

Table 5. Short Comparison with literature.

Approach	Accuracy (%)
[Rodrigues et al. 2016]	94.59
[Gual-Arnau et al. 2015]	96.10
[de Faria et al. 2018]	93.67
[da Silva et al. 2020]	92.54
[Paz-Soto et al. 2020]	98.40
[Alzubaidi et al. 2020]	99.98
Our approach	96.80

8. Conclusion

In this paper, we present an innovative method that fuses CNNs to extract features with segmented images and traditional classifiers. We validate our approach with images of red blood cells and classify them as healthy, sickle, or with other deformities. After diving into the state of the art we noticed many efforts in using CNN as a classifier neglecting the resource footprint of such approaches leading us to investigate how to improve classical classifiers.

We conducted an investigation into the performance of CNNs and conventional classifiers utilizing segmented images in our study. The achieved accuracy of 96.80% by extracting features with MobileNet and classifying with SVM demonstrates that this approach can aid in identifying sickle cell disease in microscopy images.

Although using segmented images requires a preprocessing step, the proposed approach requires less time for the classification task compared to the classification performed by the CNNs themselves while achieving similar accuracies. The proposed approach also offers a runtime approximately 14 times faster.

For future work, we intend to optimize the parameters of CNNs to further enhance their performance. Additionally, we will explore the application of other conventional classifiers to evaluate their effectiveness in our classification task. Furthermore, we plan to test the proposed methods on other datasets of medical images and exploit different CNNs architectures. These future endeavors will provide a more comprehensive understanding of the capabilities and limitations of different techniques in the field of sickle cell disease classification.

Acknowledgment

We would like to thank CAPES and FAPEMIG (Grant number CEX - APQ-02964-17) for the financial support. This study was financed in part by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001.

References

- Aljuaid, A. and Anwar, M. (2022). Survey of Supervised Learning for Medical Image Processing. *SN Computer Science*, 3(4):292.
- Alzubaidi, L., Fadhel, M. A., Al-Shamma, O., Zhang, J., and Duan, Y. (2020). Deep Learning Models for Classification of Red Blood Cells in Microscopy Images to Aid in Sickle Cell Anemia Diagnosis. *Electronics*, 9(3):427.
- Backes, A. R. (2022). Pap-smear image classification by using a fusion of texture features. In *2022 35th SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI)*, volume 1, pages 139–144.
- Carmo, N. and Mari, J. (2021). A comparative study of convolutional neural networks for classification of pigmented skin lesions. In *Anais do XVII Workshop de Visão Computacional*, pages 171–176, Porto Alegre, RS, Brasil. SBC.
- Chollet, F. et al. (2015). Keras. <https://keras.io>.
- Cortes, C. and Vapnik, V. (1995). Support-vector networks. *Machine learning*, 20(3):273–297.
- da Silva, M., Rodrigues, L., and Mari, J. F. (2020). Optimizing data augmentation policies for convolutional neural networks based on classification of sickle cells. In *Anais do XVI Workshop de Visão Computacional*, pages 46–51, Porto Alegre, RS, Brasil. SBC.
- de Faria, L. C., Rodrigues, L. F., and Mari, J. F. (2018). Cell classification using hand-crafted features and bag of visual words. In *XIV Workshop de Visão Computacional*, Ilhéus, BA, Brazil. WVC.
- Duda, R. O., Hart, P. E., and Stork, D. G. (2000). *Pattern Classification (2Nd Edition)*. Wiley-Interscience, New York, NY, USA.
- Gonzalez-Hidalgo, M., Guerrero-Pena, F., Herold-Garcia, S., Jaume-i Capó, A., and Marrero-Fernández, P. D. (2014). Red blood cell cluster separation from digital images for use in sickle cell disease. *IEEE journal of biomedical and health informatics*, 19(4):1514–1525.
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). Deep Learning. Book in preparation for MIT Press.
- Gual-Arnau, X., Herold-García, S., and Simó, A. (2015). Erythrocyte shape classification using integral-geometry-based methods. *Medical & biological engineering & computing*, 53(7):623–633.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778.

- Hoffman, R., Benz, E., Silberstein, L., Heslop, H., Anastasi, J., and Weitz, J. (2013). *Hematology: Basic Principles and Practice*. Churchill Livingstone. Saunders/Elsevier.
- Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., and Adam, H. (2017). MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. *CoRR*, abs/1704.04861.
- Huang, G., Liu, Z., and Weinberger, K. Q. (2016). Densely Connected Convolutional Networks. *CoRR*, abs/1608.06993.
- 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. (2018). Sickle cell disease. *Nature Reviews Disease Primers*, 4(1):1–22.
- Kavanagh, P. L., Fasipe, T. A., and Wun, T. (2022). Sickle Cell Disease: A Review. *JAMA*, 328(1):57–68.
- Krizhevsky, A., Sutskever, I., and E. Hinton, G. (2012). Imagenet classification with deep convolutional neural networks. *Neural Information Processing Systems*, 25.
- Nazir, S., Dickson, D. M., and Akram, M. U. (2023). Survey of explainable artificial intelligence techniques for biomedical imaging with deep neural networks. *Computers in Biology and Medicine*, 156:106668.
- Paz-Soto, Y., Herold-Garcia, S., Fernandes, L. A. F., and Díaz-Matos, S. (2020). Automatic Classification of Erythrocytes Using Artificial Neural Networks and Integral Geometry-Based Functions. In *2020 33rd SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI)*, pages 156–163, Porto de Galinhas, Brazil. IEEE.
- Petrović, N., Moyà-Alcover, G., i Capó, A. J., and González-Hidalgo, M. (2020). Sickle-cell disease diagnosis support selecting the most appropriate machine learning method: Towards a general and interpretable approach for cell morphology analysis from microscopy images. *Computers in Biology and Medicine*, 126:104027.
- Ponti, M. A., Ribeiro, L. S. F., Nazare, T. S., Bui, T., and Collomosse, J. (2017). Everything You Wanted to Know about Deep Learning for Computer Vision but Were Afraid to Ask. In *2017 30th SIBGRAPI Conference on Graphics, Patterns and Images Tutorials (SIBGRAPI-T)*, pages 17–41, Niterói, Brazil. IEEE.
- Rodrigues, L. F., Naldi, M. C., and Mari, J. F. (2016). Morphological analysis and classification of erythrocytes in microscopy images. In *Anais do XII Workshop de Visão Computacional*, Campo Grande, MS, Brazil. WVC.
- Rodrigues Moreira, L. F., Moreira, R., Travençolo, B. A. N., and Backes, A. R. (2023). An Artificial Intelligence-as-a-Service Architecture for deep learning model embodiment on low-cost devices: A case study of COVID-19 diagnosis. *Applied Soft Computing*, 134:110014.
- Schwalbe, N. and Wahl, B. (2020). Artificial intelligence and the future of global health. *The Lancet*, 395(10236):1579–1586.
- Wickramasinghe, I. and Kalutarage, H. (2021). Naive Bayes: applications, variations and vulnerabilities: a review of literature with code snippets for implementation. *Soft Computing*, 25(3):2277–2293.

- Zhou, S. K., Greenspan, H., Davatzikos, C., Duncan, J. S., Van Ginneken, B., Madabhushi, A., Prince, J. L., Rueckert, D., and Summers, R. M. (2021). A Review of Deep Learning in Medical Imaging: Imaging Traits, Technology Trends, Case Studies With Progress Highlights, and Future Promises. *Proceedings of the IEEE*, 109(5):820–838.
- Zhou, T., Ye, X., Lu, H., Zheng, X., Qiu, S., and Liu, Y. (2022). Dense Convolutional Network and Its Application in Medical Image Analysis. *BioMed Research International*, 2022:2384830.