© MAY 2026 | IRE Journals | Volume 9 Issue 11 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I11-1717227

Malaria Detection Using an Improved AlexNet-Based Deep Learning Model

M. PHANINDRA[1], K. VENKATA SIVA SAI SREEJA[2], P. NIKHIL[3], VINNAKOTA NAGA VENKATA SWATHI[4]

[1, 2, 3, 4]*Assistant Professor, Department of Computer Science & Engineering (AI & ML), R.V.R & J.C College of Engineering (Autonomous), Guntur, Andhra Pradesh, India*

***Abstract*—Diagnosing malaria still relies heavily on microscope examination of blood smears — which works, but is slow, hard to scale, and only as good as the person doing it. This paper describes a deep learning pipeline for binary malaria cell screening built on a modified AlexNet CNN. We used the public NIH/Kaggle malaria cell-image dataset, resized everything to 100×100 RGB patches, applied live augmentation, and trained a regularized CNN with batch normalization and dropout. Beyond just reporting accuracy, we evaluated the model with confusion matrices, ROC-AUC, precision-recall curves, calibration plots, and Grad-CAM heatmaps. We also compare the approach against a recent two-stage YOLOv4/DenseNet-121 system. This model doesn't do whole-slide detection or species identification — it's a clean, reproducible infected-vs-uninfected screener that can be understood, explained, and built on.**

***Index Terms*—Malaria Detection, Alexnet, Convolutional Neural Network, Grad-CAM, Medical Image Classification, Binary Screening.**

I. INTRODUCTION

Malaria is still a serious problem across tropical regions, and the standard way to diagnose it — examining stained blood films under a microscope — hasn't changed much. It works. But it requires trained staff, good equipment, and sustained attention. In busy labs, that combination is hard to maintain, which leads to delays and inconsistent results between reviewers.

Deep learning is a natural fit here. CNNs learn colour, texture, and shape patterns directly from images, so instead of manually engineering features, you train the network on labelled cell patches and let it figure out what parasites look like. The predictions are consistent in a way human reviewers often aren't. Most recent work in this space has moved toward two-stage systems: a detector finds infected cells in a full blood smear, then a classifier identifies the Plasmodium species. The journal paper we're building on does exactly this — comparing YOLOv4 and YOLOv5 for localization, then using DenseNet-121 for four-species classification. Our work uses that study as a reference point but zooms in on a narrower problem: classifying individual cell patches that have already been isolated.

That narrowing matters. If you're already working with single-cell images, you don't need object detection. The model can focus entirely on telling parasitized cells from uninfected ones, which keeps things simple enough to train, explain, and reproduce in a student project setting.

II. RELATED WORK

A. Earlier Automated Systems

Before deep learning, malaria detection pipelines used segmentation, colour thresholds, shape descriptors, and classifiers like SVM and KNN. These worked to a point, but handcrafted features broke easily — different staining protocols, lighting conditions, or microscope settings could throw everything off.

B. Deep Learning Approaches

CNNs changed that by learning feature hierarchies directly from the data. AlexNet, VGG, ResNet, MobileNet, and DenseNet have all been applied to malaria image classification with good results. YOLO-based detectors came in when the task required finding infected cells inside full smear images before classifying them.

C. What's Missing

A lot of published work either stops at binary classification or jumps straight to detection-plus-species pipelines. For a dataset of already-cropped cell images, neither extreme is ideal. The gap here is a binary classifier that's reproducible, evaluated properly, and doesn't just report one accuracy number.

© MAY 2026 | IRE Journals | Volume 9 Issue 11 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I11-1717227

III. METHODOLOGY

A. Dataset and Preprocessing

We used the Kaggle "Cell Images for Detecting Malaria" dataset from NIH. It has two folders: Parasitized and Uninfected. Images are loaded with OpenCV, converted to RGB, resized to 100×100, and stored as NumPy arrays. Class 0 is parasitized, Class 1 is uninfected.

Samples are shuffled before splitting to avoid folder-order bias. The split is manual: 25,000 training images, 2,500 for testing, the rest for validation. Class counts are printed after splitting so you can check balance before training starts.

B. Augmentation

Augmentation runs on-the-fly through Keras ImageDataGenerator — rotation up to 30 degrees, zoom, horizontal and vertical shifts, and horizontal flipping. These are reasonable transformations for cell imagery since parasite evidence doesn't disappear under moderate position or orientation changes.

C. Model Architecture

The model is based on AlexNet but adapted for 100×100 patches. It has five convolutional layers, max-pooling blocks, batch normalization after the major conv stages, two dense layers of 512 neurons, and dropout at 0.5. The output is a single sigmoid neuron for binary prediction. Training uses Adam and binary cross-entropy over 12 epochs with batch size 64.

Component	Configuration
Input size	100 x 100 x 3
Classes	Parasitized, Uninfected
Optimizer	Adam
Loss	Binary cross-entropy
Regularization	Batch normalization, dropout
Evaluation	ROC, PR, calibration, Grad-CAM

Table I: Summary of the proposed AlexNet-based pipeline.

IV. RESULTS AND DISCUSSIONS

The model performs well on binary classification. ROC-AUC and average precision both come in around 0.98, which means it separates the two classes well across decision thresholds — not just at one particular cutoff. Confusion matrices and classification reports break down false positives and false negatives more concretely.

Calibration analysis is worth including in medical screening work. A model can rank images correctly and still produce confidence scores that mean nothing in practice. Checking calibration at least tells you whether the probabilities are interpretable. Grad-CAM heatmaps go a step further — they show which regions of the image are actually driving predictions, so you can check whether the network is looking at parasite-related structures or just noise.

Metric/View	Purpose
Accuracy	Overall correctness
ROC-AUC	Threshold-independent separation
Average precision	Precision-recall quality
Sensitivity	Ability to detect infected cells
Specificity	Ability to reject uninfected cells
Grad-CAM	Visual explanation of model focus

Table II: Evaluation views used in the implementation.

To be clear about the comparison: the journal framework this work references is solving a harder problem. It handles whole-image localization and four-species classification, with DenseNet-121 hitting 95.5% species accuracy. This model doesn't compete with that — it's a binary screener for pre-cropped patches, not a full diagnostic pipeline.

V. LIMITATIONS AND FUTURE WORK

The model doesn't do full-smear detection, bounding-box localization, or species classification. The manual train/test split could be improved with stratified partitioning, and validation augmentation should be separated from training augmentation in a cleaner experiment setup. Explicit normalization and better experiment logging would also help reproducibility.

From here, the obvious next step is comparing AlexNet against DenseNet-121, MobileNetV2, and other transfer-learning baselines on the same

dataset. After that, extending the pipeline toward whole-smear object detection and eventually multiclass species recognition would bring it closer to something clinically useful.

VI. CONCLUSION

This paper presents a malaria cell screener built on a modified AlexNet CNN. By working with already-cropped cell patches, we avoided the complexity of localization and kept the model focused on a single task it can do well. The evaluation goes well beyond accuracy — augmentation, batch normalization, dropout, ROC-AUC, precision-recall, calibration, and Grad-CAM all contribute to a more trustworthy result. Think of this as a solid binary screening foundation, not a finished diagnostic tool. The path toward that is clear enough from here.

VII. ACKNOWLEDGMENT

The authors thank Mrs. Vinnakota Naga Venkata Swathi, Assistant Professor, CSE (AI & ML), for project guidance, and Dr. G. Rama Mohan Babu, Head of Department, and Dr. Kolla Srinivas, Principal of R.V.R & J.C College of Engineering, for institutional support.

REFERENCES

[1] D. Sukumarran et al., "Automated Identification of Malaria-Infected Cells and Classification of Human Malaria Parasites Using a Two-Stage Deep Learning Technique," IEEE Access, vol. 12, 2024.
[2] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," NeurIPS, 2012.
[3] G. Huang et al., "Densely Connected Convolutional Networks," CVPR, 2017.
[4] R. R. Selvaraju et al., "Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization," ICCV, 2017.
[5] I. Arunava, "Cell Images for Detecting Malaria," Kaggle Dataset.
[6] S. Rajaraman et al., "Pre-trained convolutional neural networks as feature extractors toward improved malaria parasite detection," PeerJ, 2018.