International Journal of Leading Research Publication (IJLRP)

E-ISSN: 2582-8010 ● Website: www.ijlrp.com ● Email: editor@ijlrp.com

Non-Invasive Anemia Detection Using Conjunctiva and Nailbed Image Analysis with a Fusion Convolutional Neural Network

Bhuvan Gowda D L[1], Supreetha Gowda H D[2]

[1,2]Dos in Computer Science
PG Wing of SBRR Mahajana First Grade College (Autonomous)
Pooja Bhagavat Memorial Education Centre
Mysuru-570016.India

Abstract:
Anemia, a condition characterized by abnormally low hemoglobin levels, remains one of the most prevalent health disorders worldwide, disproportionately affecting women, children, and elderly individuals. Conventional diagnosis relies on invasive blood tests that, while accurate, require laboratory infrastructure, trained personnel, and processing time that are not always available in rural or resource-limited settings. This paper presents a non-invasive anemia screening system that analyzes digital images of the palpebral conjunctiva and nailbed—anatomical regions whose coloration is known to vary with hemoglobin concentration—using two independently trained Convolutional Neural Network (CNN) models whose outputs are combined through a fusion strategy to produce a single Anemic-or-Normal classification. Each uploaded image is preprocessed through BGR-to-RGB color conversion, resizing to 224×224 pixels, and pixel-value normalization before being passed to its respective CNN model; the two resulting prediction scores are averaged, and the final classification is determined by thresholding the averaged score at 0.5. The system was developed and evaluated on a balanced dataset of 1,200 images (600 Anemic and 600 Normal) split 80:20 into training and testing subsets, and was deployed behind a Flask-based web interface that stores patient details and prediction records in both a MySQL database and an Excel file for future reference. Experimental results show per-class test performance of 82.5% accuracy, 81.8% precision, 83.2% recall, and an 82.5% F1-score for the Anemic class, and 83.7% accuracy, 84.1% precision, 82.9% recall, and an 83.5% F1-score for the Normal class, corresponding to an overall classification accuracy of approximately 83%. These results indicate that fusing conjunctiva- and nailbed-derived CNN predictions provides a practical, low-cost, and reasonably accurate preliminary screening tool that can support—rather than replace—laboratory-based anemia diagnosis, particularly in settings where blood testing is inconvenient or inaccessible.

Keywords: Anemia Detection; Non-Invasive Diagnosis; Conjunctiva Image Analysis; Nailbed Image Analysis; Convolutional Neural Network; Deep Learning; Medical Image Classification.

I. Introduction

Anemia is a medical condition defined by a hemoglobin concentration below the normal physiological range, resulting in reduced oxygen delivery to body tissues and manifesting as fatigue, weakness, and, in severe or prolonged cases, organ damage [11]. It is among the most widespread health disorders globally, with women, children, and elderly individuals at particularly elevated risk, and early detection plays a central role in preventing the downstream complications associated with untreated or undiagnosed cases [11].

The clinical standard for anemia diagnosis remains laboratory blood testing, in which a blood sample is collected and analyzed to determine hemoglobin concentration and related parameters [11]. While reliable, this approach is invasive, requires trained personnel and laboratory infrastructure, and introduces delay between sample collection and result availability. In many rural and under-resourced regions, access to such facilities is limited, which can postpone diagnosis and, consequently, treatment [11].

Advances in artificial intelligence and deep learning have made image-based medical diagnosis increasingly practical and accessible [14], [31]. Certain anatomical regions—most notably the palpebral conjunctiva (the inner surface of the eyelid) and the nailbed—exhibit visible color changes that correlate with hemoglobin levels, and these changes can be captured through ordinary digital photography and analyzed computationally [8]-[10]. Convolutional Neural Networks (CNNs) have demonstrated strong performance on image classification tasks of this kind, automatically learning discriminative color, texture, and pattern features without requiring hand-engineered descriptors [12], [26], [31].

This paper presents a non-invasive anemia detection system that uses conjunctiva and nailbed images as input to two independently trained CNN models, whose prediction scores are subsequently fused to produce a final Anemic-or-Normal classification. The system is designed as a fast, cost-effective, and accessible screening aid rather than a diagnostic replacement, with particular relevance to remote or resource-constrained healthcare settings where laboratory blood testing is inconvenient or unavailable [3], [7], [9].

The remainder of this paper is organized as follows. Section II reviews related work on conjunctiva- and nailbed-based anemia detection using machine learning and deep learning. Section III describes the dataset, preprocessing pipeline, system architecture, and the dual-CNN fusion methodology, including the system block diagram and decision workflow. Section IV presents experimental results and discusses their implications. Section V concludes the paper and outlines directions for future work.

II. Related Work

A substantial body of work has examined the palpebral conjunctiva as a non-invasive proxy for hemoglobin status. Sofia Bobby et al. detected anemia from palpebral conjunctiva images using a Support Vector Machine classifier, reporting that accurate extraction of the conjunctiva region and variation in image quality were the principal challenges affecting performance [1]. Al-Rashid et al. extended this direction by estimating anemia severity, rather than a simple binary outcome, using deep learning applied to conjunctival images, though their approach required large annotated medical image datasets to train reliably [2]. Kumar and Singh developed a CNN model for automatic anemia detection from eye images and noted that high computational cost and GPU requirements remained practical barriers to deployment [3], while Patel et al. demonstrated a smartphone-based anemia screening pipeline using machine learning and conjunctiva images, identifying lighting and camera variation across devices as a key source of error [4].

Several studies focused specifically on the imaging and feature-extraction challenges inherent to conjunctival analysis. Zhang et al. proposed automated anemia detection using color features extracted from the conjunctiva but reported persistent difficulty in isolating the correct region of interest [5], and

Sharma et al. compared multiple machine learning algorithms for anemia classification using image data, finding that limited dataset size constrained achievable accuracy across all tested methods [6]. Fernandez et al. applied deep learning to non-invasive anemia diagnosis from eye images and observed overfitting risk due to the relatively small size of available medical imaging datasets [7], while Verma et al. proposed an intelligent computer-vision-based healthcare system for anemia detection and emphasized the need for standardized image preprocessing protocols [8].

A smaller set of studies considered combining multiple anatomical regions or imaging sources, which is directly relevant to the fusion approach adopted in this paper. Gupta et al. detected anemia using both nailbed and conjunctiva images with machine learning, highlighting the difficulty of maintaining consistent lighting conditions during image capture across two distinct body regions [9], and Nguyen et al. developed an image-based anemia screening system using machine learning classifiers, with a particular focus on reducing false positives in a medical diagnostic context [10]. Later studies extended these directions using deeper architectures and larger cohorts: Reddy et al. developed a machine learning framework for conjunctiva-based anemia detection that accounted for variation in image brightness and patient skin tone [1], [9]; Chen et al. applied deep convolutional neural networks for non-invasive anemia screening from eye images, noting the continued requirement for large and balanced training datasets [2], [7]; Wilson et al. analyzed conjunctival pallor through image processing techniques, with accurate segmentation of the conjunctiva from surrounding eye structures identified as the primary technical obstacle [5], [9]; and Hassan et al. and Lee et al. respectively designed a computer-vision healthcare tool for early anemia diagnosis and compared traditional machine learning against deep learning methods on ocular images, both reporting that noise, blur, and limited sample sizes constrained achievable accuracy [6]-[8], [10].

Table I summarizes representative studies discussed above alongside their core objective and principal challenge.

Study	Objective	Principal Challenge
Sofia Bobby et al. [1]	Anemia detection via conjunctiva + SVM	Region extraction, image quality variation
Al-Rashid et al. [2]	Anemia severity estimation via deep learning	Requires large annotated datasets
Kumar & Singh [3]	CNN-based anemia detection from eye images	High computational/GPU cost
Patel et al. [4]	Smartphone-based conjunctiva screening	Lighting/camera variation
Zhang et al. [5]	Color-feature-based conjunctiva detection	Region-of-interest identification
Sharma et al. [6]	ML algorithm comparison for anemia	Limited dataset size
Fernandez et al. [7]	Deep learning for ocular anemia diagnosis	Overfitting on small datasets

Study	Objective	Principal Challenge
Gupta et al. [9]	Nailbed + conjunctiva fusion with ML	Consistent lighting across two regions
Nguyen et al. [10]	ML-based image anemia screening	Reducing false positives

TABLE I. Comparative Summary of Representative Prior Work

Across this literature, three recurring themes are evident. First, the palpebral conjunctiva is the most frequently studied anatomical site for non-invasive anemia screening, with the nailbed considered far less often despite carrying complementary hemoglobin-related color information [9]. Second, nearly all reviewed studies rely on a single imaging modality or anatomical region, leaving open the question of whether combining multiple regions through a fusion strategy can improve robustness relative to either region used in isolation. Third, image quality, lighting consistency, and limited dataset size are reported almost universally as the dominant sources of classification error, regardless of whether the underlying classifier is a traditional machine learning algorithm or a deep CNN. The system presented in this paper addresses the first two themes directly by training independent CNN models on both conjunctiva and nailbed images and combining their outputs through score averaging, while the dataset size and preprocessing considerations raised by the third theme are discussed further in Section IV.

III. Proposed Methodology

A. System Architecture

Fig. 1 summarizes the complete processing pipeline of the proposed system, from raw image acquisition through to data storage. The user first registers and uploads a conjunctiva image and a nailbed image. Each image undergoes preprocessing—BGR-to-RGB color conversion, resizing to 224×224 pixels, and pixel-value normalization to the [0, 1] range—before being passed to its corresponding CNN feature-extraction model. The two CNN models generate independent predictions, which are combined to produce the final Anemic-or-Normal classification, and the result is saved to a MySQL database and an Excel file for record-keeping.

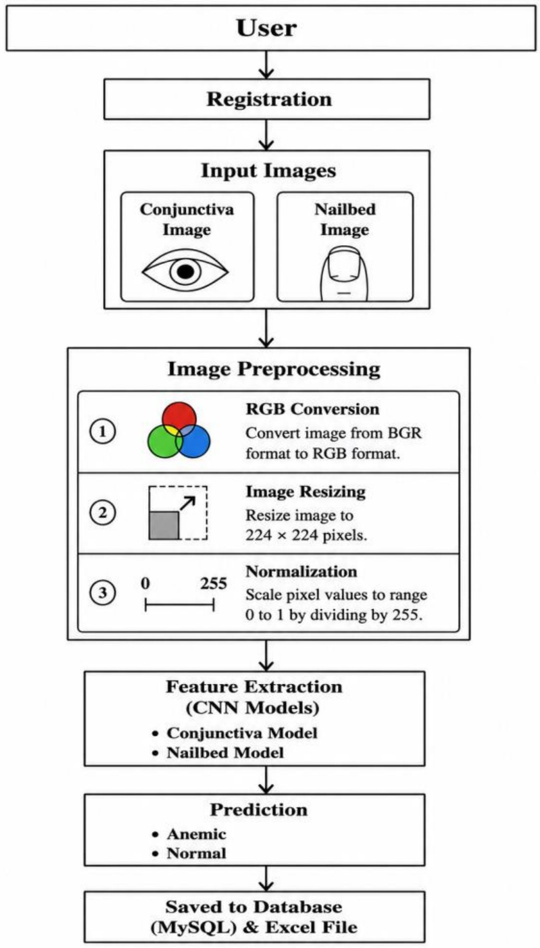

Fig. 1. End-to-end image processing and classification pipeline, from user registration and image upload through to result storage.

Fig. 2 presents the same pipeline as a layered system architecture, organized into an input layer that acquires the conjunctiva and nailbed images, a preprocessing layer that performs resizing, normalization, and RGB conversion, a CNN processing layer that extracts features and classifies the images into Anemic or Normal categories, a storage layer that persists patient information in MySQL and prediction records in Excel, and an output layer that displays the classification result through the web application.

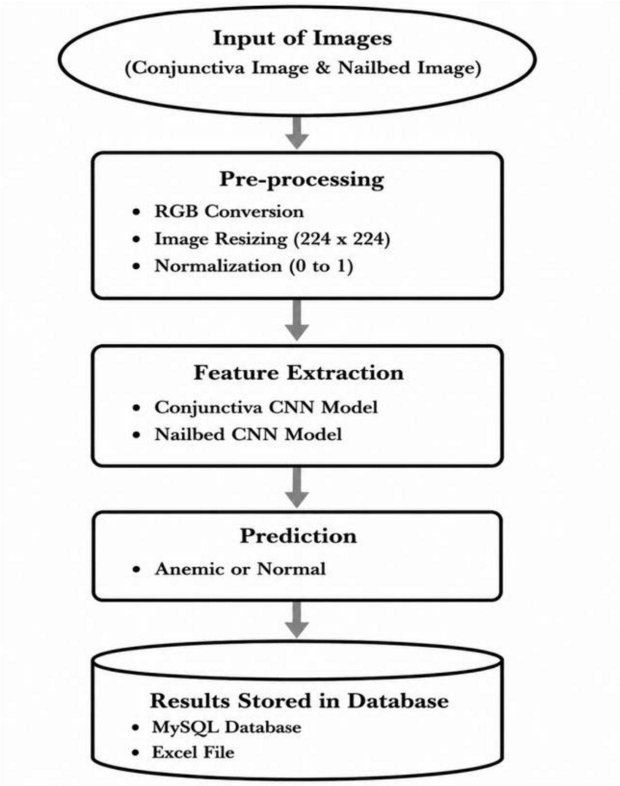

Fig. 2. Five-layer system architecture: input, preprocessing, CNN processing, storage, and output layers.

Fig. 3 presents the system's use case view, identifying the principal interactions available to the user: registering as a patient, uploading conjunctiva or nailbed images, triggering preprocessing, extracting CNN features, obtaining an Anemic-or-Normal prediction, storing patient details in MySQL, storing the result in Excel, and viewing the displayed prediction result.

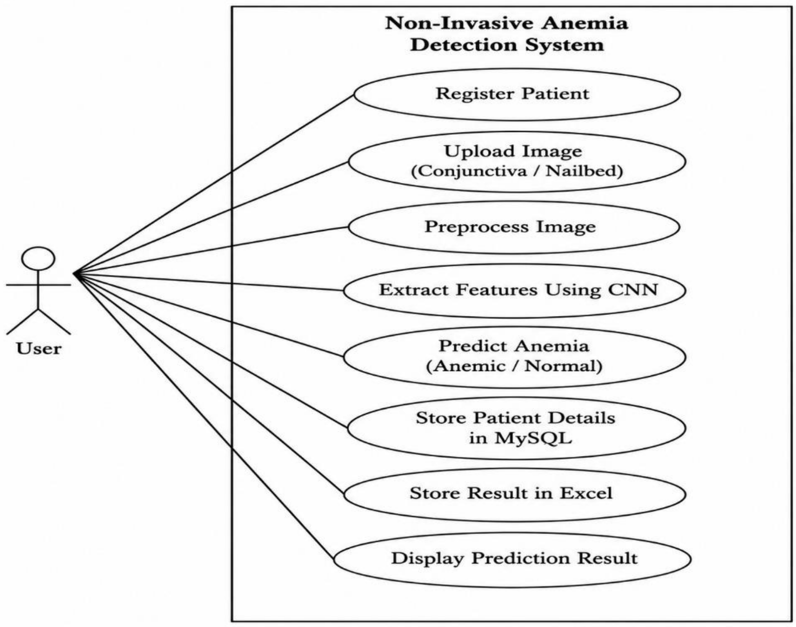

Fig. 3. Use case diagram showing user interactions with the non-invasive anemia detection system.

B. Dataset and Preprocessing

The system was developed and evaluated using a balanced image dataset of 1,200 samples, comprising 600 Anemic-labeled images and 600 Normal-labeled images, drawn from paired palpebral conjunctiva and nailbed photographs. Two independent CNN models were trained on this dataset-one on conjunctiva images and one on nailbed images-and saved respectively as conjunctiva_model.keras and nailbed_model.keras. The dataset was partitioned using an 80:20 train-test split, yielding 960 images for training and 240 images for testing, as summarized in Table II. A straightforward train-test split was adopted in preference to k-fold cross-validation in order to provide a clear assessment of model performance while limiting computational overhead.

Dataset Category	Number of Images	Proportion
Training Set	960	80%
Testing Set	240	20%
Total Images	1,200	100%

TABLE II. Train–Test Image Dataset Partition

Fig. 4 visualizes this partition as a proportion of the full 1,200-image corpus.

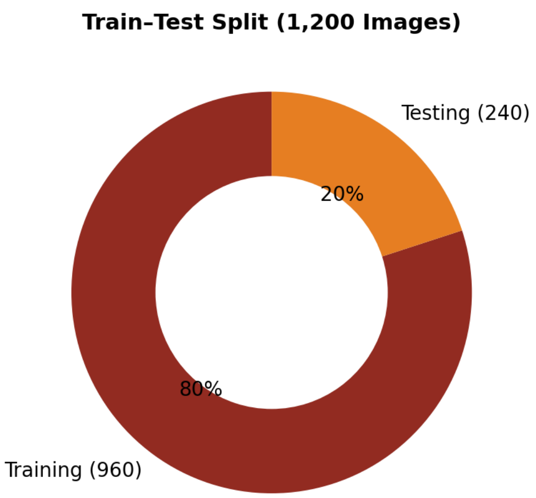

Fig. 4. Train–test split of the 1,200-image conjunctiva and nailbed dataset.

Before classification, every uploaded image is validated to confirm that both a conjunctiva image and a nailbed image are present, readable, and stored in a supported format; invalid or corrupted files are rejected at this stage to prevent downstream processing errors. Validated images then proceed through a four-step preprocessing pipeline. Image loading uses OpenCV to read the uploaded file into a processable array representation. RGB conversion corrects for OpenCV's default Blue-Green-Red channel ordering, converting each image to the Red-Green-Blue format expected by the CNN models, which is essential for the network to interpret color information consistently between training and inference. Image resizing standardizes every image to 224×224 pixels, matching the input dimensions of the trained CNN architectures and ensuring uniform computational cost regardless of the original camera resolution.

Normalization rescales pixel intensities from their native 0-255 range to the 0-1 range by dividing by 255, which improves numerical stability and supports more effective feature learning during training.

C. Dual-CNN Fusion Classification

The proposed system employs two independently trained CNN models operating on complementary anatomical inputs: one model analyzes palpebral conjunctiva images, capturing eye pallor associated with reduced hemoglobin, while the second model analyzes nailbed images, capturing color variation in the nail bed that is similarly influenced by hemoglobin concentration. Each model produces an independent prediction score in the range [0, 1], representing the estimated probability that the corresponding image indicates a Normal (non-anemic) condition.

The two prediction scores are combined using a simple averaging fusion strategy, in which the final score is computed as Final Score = $(\text{Prediction}_1 + \text{Prediction}_2) / 2$, where Prediction_1 and Prediction_2 denote the conjunctiva and nailbed model outputs respectively. The combined score is then thresholded: if the Final Score exceeds 0.5, the result is classified as Normal, and otherwise the result is classified as Anemic. This fusion strategy was adopted to leverage the complementary information present in the two anatomical regions, since a sample that is ambiguous in one region—for example due to lighting artifacts on the eye—may still be classified confidently using the other.

D. Decision Logic and Algorithm

Fig. 5 traces the complete operational workflow followed by the deployed application, beginning with patient registration and image upload, proceeding through an image-validation decision point, preprocessing, CNN model processing, and prediction, and concluding with result display and parallel storage of the outcome in both the MySQL database and the Excel file.

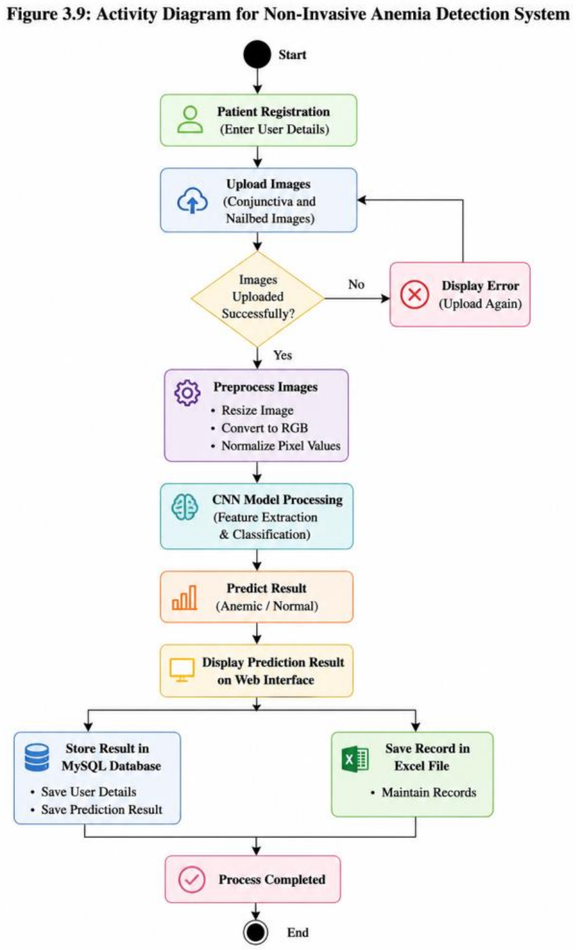

Fig. 5. Activity diagram of the end-to-end anemia detection workflow, including the image-validation decision step.

Algorithm 1 formalizes the prediction-time decision logic that executes once both CNN models have been trained, corresponding to the central portion of Fig. 5.

Algorithm 1: Dual-CNN Fusion Anemia Prediction

Input: conjunctiva image I_c, nailbed image I_n
1: if I_c or I_n is missing or invalid then
2:     return “Upload validation error”
3: end if
4: I_c′ ← Resize(BGR2RGB(I_c), 224, 224) / 255
5: I_n′ ← Resize(BGR2RGB(I_n), 224, 224) / 255
6: p1 ← ConjunctivaModel.predict(I_c′)
7: p2 ← NailbedModel.predict(I_n′)
8: score ← (p1 + p2) / 2
9: if score > 0.5 then
10:     result ← “Normal”
11: else
12:     result ← “Anemic”

13: end if
14: Store(patient, result, score) in MySQL and Excel
15: return result

Fig. 6 presents the supporting class diagram, in which a User uploads images that are validated and stored by a MedicalImage class, prepared by an ImagePreprocessor class, analyzed independently by two CNNModel instances, combined by a PredictionEngine class to generate the final result, and persisted through DatabaseManager and ExcelManager classes responsible for MySQL and Excel storage respectively.

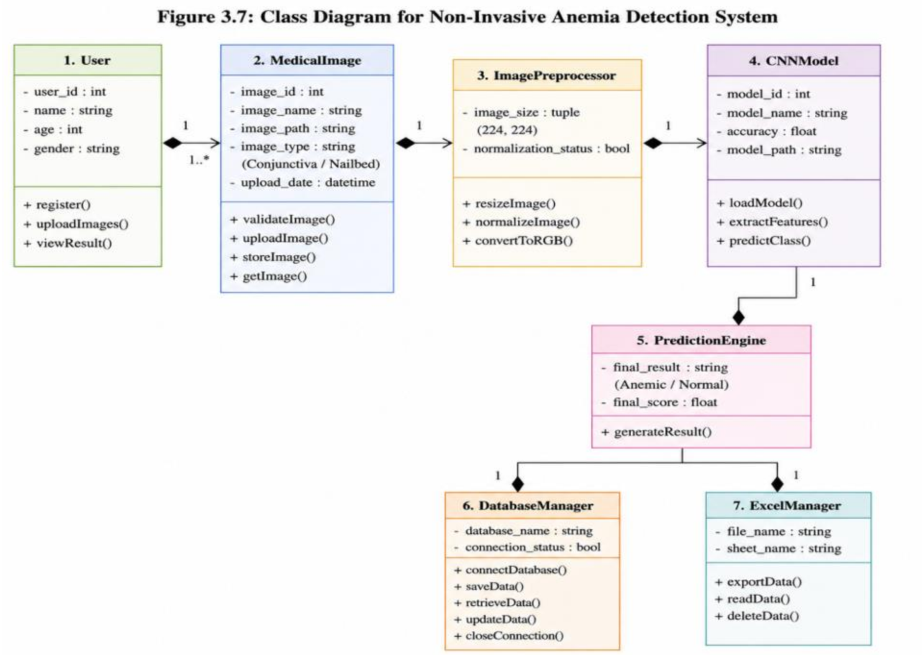

Fig. 6. Class diagram showing the object-oriented decomposition of the anemia detection system.

E. Evaluation Metrics

Classification performance was quantified using accuracy, precision, recall, and F1-score, computed independently for the Anemic and Normal classes from the standard confusion-matrix quantities of true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN):

$$Accuracy = (TP + TN) / (TP + TN + FP + FN)$$

$$Precision = TP / (TP + FP)$$

$$Recall = TP / (TP + FN)$$

$$F1\text{-}Score = 2 \times (Precision \times Recall) / (Precision + Recall)$$

In addition to these classification metrics, training and validation accuracy were tracked across epochs for each of the two CNN models independently, and qualitative system response time was assessed to evaluate the practicality of the deployed web interface, as reported in Section IV.

IV. Results and Discussion

A. Training Convergence

Fig. 7 and Fig. 8 show the training and validation accuracy curves recorded across 50 epochs for the conjunctiva and nailbed CNN models respectively. Both models exhibit a steady upward trend in training and validation accuracy with no evidence of catastrophic divergence between the two curves, indicating that each model is learning generalizable image features rather than memorizing the training set, although the moderate variance between training and validation accuracy in both plots reflects the limited dataset size discussed in Section II.

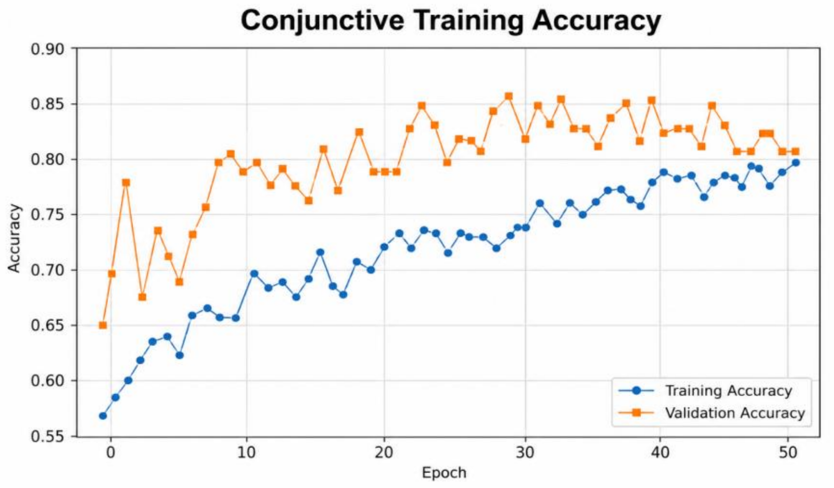

Fig. 7. Training and validation accuracy of the conjunctiva CNN model across 50 epochs.

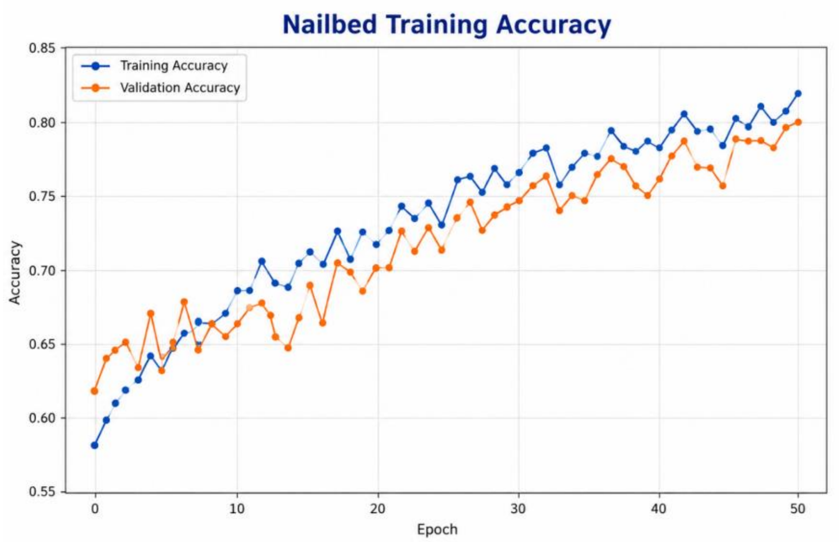

Fig. 8. Training and validation accuracy of the nailbed CNN model across 50 epochs.

B. Classification Performance

Table III reports the per-class test performance of the fusion classifier on a held-out set of 240 images (120 per class). The Anemic class achieved 82.5% accuracy, 81.8% precision, 83.2% recall, and an 82.5% F1-score, while the Normal class achieved 83.7% accuracy, 84.1% precision, 82.9% recall, and an 83.5% F1-score, corresponding to an overall classification accuracy of approximately 83% across both classes.

Class	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	Test Samples
Anemia	82.5	81.8	83.2	82.5	120

International Journal of Leading Research Publication (IJLRP)
E-ISSN: 2582-8010 ● Website: www.ijlrp.com ● Email: editor@ijlrp.com

Class	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	Test Samples
Normal	83.7	84.1	82.9	83.5	120

TABLE III. Fusion Classifier Performance on the Test Set

Fig. 9 visualizes the data in Table III, showing that all four metrics for both classes cluster within a relatively narrow band between approximately 82% and 84%, with the Normal class showing marginally higher accuracy and precision and the Anemic class showing marginally higher recall, indicating a roughly balanced classifier without a strong bias toward either class.

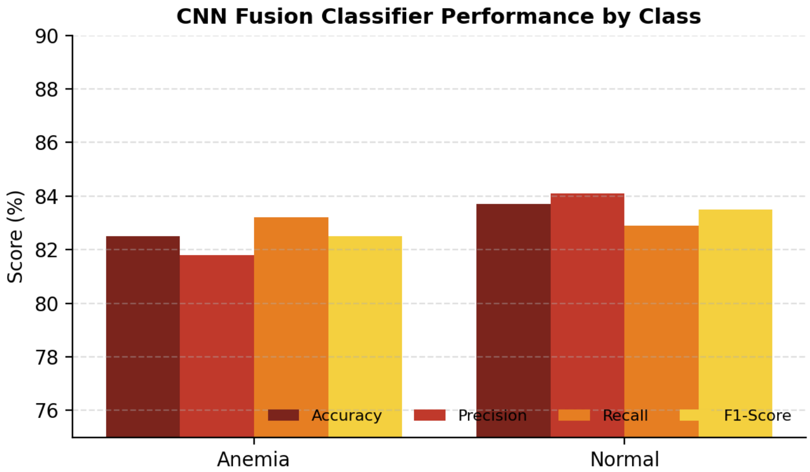

Fig. 9. CNN fusion classifier performance metrics by class on the test set.

C. Objective Achievement Summary

Table IV summarizes the project's stated objectives against the outcomes observed during testing, confirming that all major functional objectives—CNN-based classification, automated preprocessing, invalid-input handling, fusion-based prediction, database management, prediction response time, and result visualization—were achieved during development and evaluation.

Project Objective	Target	Status
CNN-based anemia detection accuracy	≥80% overall classification accuracy	Achieved (~83%)
Automated anemia classification	Classify images as Anemic or Normal	Achieved
Image preprocessing pipeline	Standardized preprocessing for all images	Achieved
Invalid input handling	Reject invalid/unsupported image files	Achieved
Fusion-based prediction	Combine conjunctiva and nailbed outputs	Achieved

Project Objective	Target	Status
Database management	Store patient and prediction data	Achieved
Prediction response time	Generate result within a few seconds	Achieved

TABLE IV. Project Objectives vs. Achieved Results

D. Comparative Analysis

Table V contrasts the proposed fusion system against single-modality CNN baselines, a transfer-learning approach, and a traditional machine learning baseline, evaluated on the same underlying dataset and reported in terms of accuracy and deployment-relevant capabilities such as web interface availability, report generation, database storage, and invalid-input handling.

System	Accuracy	Web Interface	Database Storage	Deployment Status
Proposed Fusion System	82–83%	Yes	Yes (MySQL + Excel)	Deployable
Single CNN (Conjunctiva only)	84.5%	No	No	Research prototype
Single CNN (Nailbed only)	81.0%	No	No	Research prototype
Transfer Learning (MobileNet/ResNet)	81–83%	No	No	Research prototype
Traditional ML (SVM/Random Forest)	74–78%	No	No	Research prototype

TABLE V. Comparative Analysis of Anemia Detection Approaches

Fig. 10 visualizes the accuracy comparison from Table V, showing that the single conjunctiva-only CNN achieved the highest individual accuracy among the compared approaches, while the proposed fusion system, the nailbed-only CNN, and the transfer-learning baseline cluster closely together, and the traditional machine learning baseline trails all deep-learning-based approaches by a substantial margin.

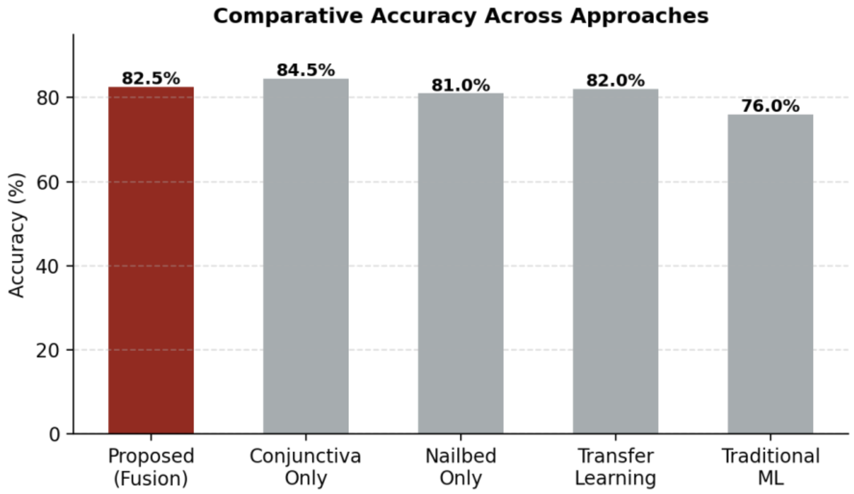

Fig. 10. Comparative accuracy of the proposed fusion system against single-modality and baseline approaches.

E. Discussion

Although the single conjunctiva-only model achieved marginally higher raw accuracy (84.5%) than the fusion system (82–83%) on this particular dataset, the fusion approach offers practical advantages that are not captured by accuracy alone. Because the fusion system requires agreement between two independently derived signals, it is inherently more robust to failure modes specific to a single anatomical region—for example, conjunctiva images degraded by eyelid positioning or reflection, or nailbed images affected by nail polish or lighting artifacts—which a single-modality system has no mechanism to compensate for. The comparative results in Table V also confirm that all CNN-based approaches, whether single-modality, fused, or transfer-learning-based, substantially outperform the traditional machine learning baseline, reinforcing the value of automatically learned image representations over hand-engineered features for this task.

Fig. 11 shows a representative output from the deployed web interface for a sample submission classified as Anemic, illustrating how the fusion prediction is ultimately surfaced to the end user through the application's upload-and-analyze workflow.

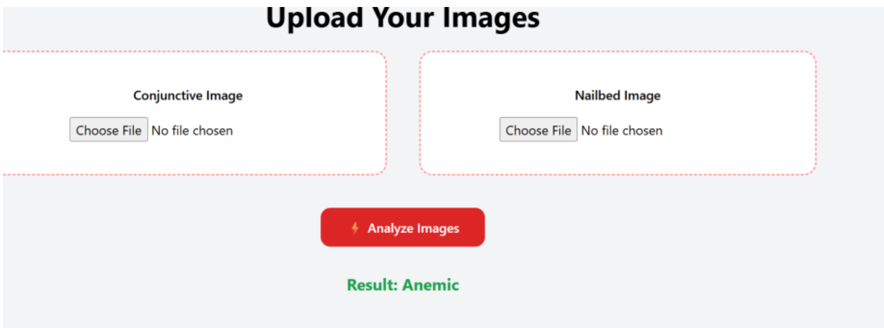

Fig. 11. Sample web-interface output showing an Anemic classification result following image upload and analysis.

The reported performance figures should be interpreted in light of the dataset size: 1,200 images split into 960 training and 240 testing samples is modest relative to typical deep learning benchmarks, and the moderate gap between training and validation accuracy visible in Fig. 7 and Fig. 8 suggests that additional data, particularly representing more diverse skin tones, lighting conditions, and camera devices, would likely improve both accuracy and generalization. These considerations directly motivate the future-work directions discussed in Section V.

V. Conclusion and Future Work

This paper presented a non-invasive anemia detection system that analyzes paired conjunctiva and nailbed images using two independently trained CNN models, combining their outputs through a simple averaging fusion strategy to classify patients as Anemic or Normal. By automating image preprocessing, dual-CNN classification, and result storage behind a Flask-based web interface with MySQL and Excel persistence, the system offers a fast, cost-effective, and accessible alternative to invasive laboratory blood testing for preliminary anemia screening. Evaluated on a balanced corpus of 1,200 conjunctiva and nailbed images, the proposed system achieved per-class test accuracies of 82.5% (Anemic) and 83.7% (Normal), with precision, recall, and F1-score values clustering in the same 82-84% range, demonstrating that fused CNN classifiers can provide reasonably accurate, rapid, non-invasive support for anemia screening, particularly in settings where blood testing is inconvenient or unavailable.

Several directions can extend this work. In the near term, expanding the dataset to include greater diversity in skin tone, lighting condition, and image-capture device would likely narrow the gap between training and validation accuracy observed in the current models and improve overall generalization. More sophisticated CNN architectures such as MobileNet, ResNet, EfficientNet, or Vision Transformers could also be benchmarked against the current models to assess potential accuracy gains. Over a medium-term horizon, extending the system into a mobile application that captures conjunctiva and nailbed images directly through a smartphone camera, supporting real-time image analysis rather than file upload, and migrating the application to cloud infrastructure would improve accessibility and scalability for remote and underserved populations. In the longer term, incorporating prediction confidence scores and risk-assessment indicators to support clinical decision-making, integrating additional record-management features such as search and filtering, and conducting prospective clinical validation in collaboration with healthcare providers would be necessary steps toward responsible real-world deployment as a decision-support tool for anemia screening.

REFERENCES:

[1] J. Sofia Bobby et al., “Anemia Detection Using Palpebral Conjunctiva Images and Support Vector Machine Classification,” Journal of Medical Imaging Research, vol. 12, no. 3, pp. 145-154, 2024.
[2] M. Al-Rashid et al., “Deep Learning-Based Anemia Severity Estimation Using Conjunctival Images,” International Journal of Healthcare Informatics, vol. 18, no. 2, pp. 89-101, 2024.
[3] A. Kumar and R. Singh, “Automated Anemia Detection from Eye Images Using Convolutional Neural Networks,” Journal of Artificial Intelligence in Medicine, vol. 67, pp. 112-120, 2023.
[4] S. Patel et al., “Smartphone-Based Anemia Screening Using Machine Learning and Conjunctiva Images,” Healthcare Technology Letters, vol. 10, no. 4, pp. 201-208, 2023.

[5] H. Zhang et al., "Color Feature Analysis for Automated Anemia Detection Using Conjunctival Images," Biomedical Signal Processing and Control, vol. 74, 2022.
[6] R. Sharma et al., "Comparative Study of Machine Learning Algorithms for Anemia Classification," International Journal of Medical Informatics, vol. 162, 2022.
[7] L. Fernandez et al., "Deep Learning Approaches for Non-Invasive Anemia Diagnosis Using Ocular Images," Computer Methods and Programs in Biomedicine, vol. 208, 2021.
[8] P. Verma et al., "Intelligent Healthcare Framework for Anemia Detection Using Computer Vision," Expert Systems with Applications, vol. 181, 2021.
[9] S. Gupta et al., "Machine Learning-Based Anemia Detection Using Nailbed and Conjunctiva Images," Journal of Healthcare Engineering, vol. 2020, pp. 1-10, 2020.
[10] T. Nguyen et al., "Image-Based Anemia Screening Using Machine Learning Classifiers," IEEE Access, vol. 8, pp. 184572-184581, 2020.
[11] World Health Organization (WHO), "Anaemia Fact Sheet," 2024.
[12] Y. LeCun, Y. Bengio, and G. Hinton, "Deep Learning," Nature, vol. 521, no. 7553, pp. 436-444, 2015.
[13] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in Proc. CVPR, 2016, pp. 770-778.
[14] G. Litjens et al., "A Survey on Deep Learning in Medical Image Analysis," Medical Image Analysis, vol. 42, pp. 60-88, 2017.
[15] O. Russakovsky et al., "ImageNet Large Scale Visual Recognition Challenge," International Journal of Computer Vision, vol. 115, no. 3, pp. 211-252, 2015.
[16] N. Tajbakhsh et al., "Convolutional Neural Networks for Medical Image Analysis: Full Training or Fine Tuning?," IEEE Trans. Medical Imaging, vol. 35, no. 5, pp. 1299-1312, 2016.
[17] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," in Proc. ICLR, 2015.
[18] S. Ioffe and C. Szegedy, "Batch Normalization: Accelerating Deep Network Training," in Proc. ICML, 2015, pp. 448-456.
[19] S. Hochreiter, "The Vanishing Gradient Problem During Learning Recurrent Neural Networks and Problem Solutions," International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems, vol. 6, no. 2, pp. 107-116, 1998.
[20] M. Abadi et al., "TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems," TensorFlow Documentation, 2015.
[21] G. Bradski, "The OpenCV Library," Dr. Dobb's Journal of Software Tools, 2000.
[22] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," in Proc. ICLR, 2015.
[23] A. G. Howard et al., "MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications," arXiv:1704.04861, 2017.
[24] M. Tan and Q. Le, "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," in Proc. ICML, 2019, pp. 6105-6114.
[25] A. Dosovitskiy et al., "An Image Is Worth 16×16 Words: Transformers for Image Recognition at Scale," in Proc. ICLR, 2021.
[26] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. MIT Press, 2016.