

# Deep Learning Based Non-Invasive Framework for Nutritional Deficiency Detection Using Hair and Nail Images

B. Shilpa Shree<sup>1,\*</sup>, D. Sumathi<sup>2</sup>, Prem Kumar Ramesh<sup>3</sup>, M. Sindhu<sup>4</sup>, S. P. Indushri<sup>5</sup>, S. Naga Lahari<sup>6</sup>, Farrukh Arslan<sup>7</sup>

<sup>1,2,3,4,5,6</sup>Department of Computer Science and Engineering, CMR Institute of Technology, Bengaluru, Karnataka, India.

<sup>7</sup>Department of Computer Science, Purdue University, West Lafayette, Indiana, United States of America.  
shilpashree.b@cmrit.ac.in<sup>1</sup>, d.sumathiphd@gmail.com<sup>2</sup>, premkumar.r@cmrit.ac.in<sup>3</sup>, sim22cs@cmrit.ac.in<sup>4</sup>,  
spin22cs@cmrit.ac.in<sup>5</sup>, snla22cs@cmrit.ac.in<sup>6</sup>, farslan@purdue.edu<sup>7</sup>

**Abstract:** Nutritional inadequacies are a big worry these days. Iron, zinc, biotin, and other vitamin deficiencies can cause a range of health problems, including anemia, a weakened immune system, hair loss, and nails that aren't growing properly. The problems that were talked about above affect everyone, but they are especially bad for women, children, and those who come from less fortunate socioeconomic backgrounds. The most common way to detect deficiencies is through expensive blood tests, which makes it hard for people to get tested regularly. This research proposes a pragmatic methodology using Convolutional Neural Networks (CNNs) to precisely detect deficiencies in hair and nails by analyzing high-resolution images. Our method uses a CNN trained on images of hair and nails. It has an 89% accuracy rate in distinguishing between different types of deficiencies, including iron, zinc, biotin, and vitamin deficiencies. The model uses several convolutional layers to process images of size  $224 \times 224$  pixels. It also uses data augmentation to improve its accuracy. This technique also helps people in communities that aren't getting enough aid get an early diagnosis at a low cost.

**Keywords:** Convolutional Neural Networks (CNNs); Nutritional Deficiency; Hair and Nail Analysis; Medical Imaging; Healthcare Screening; Vitamin Detection; Traditional Diagnosis; Health Problems.

**Received on:** 03/04/2025, **Revised on:** 08/06/2025, **Accepted on:** 01/09/2025, **Published on:** 08/03/2026

**Journal Homepage:** <https://www.fmdbpub.com/user/journals/details/FTSHSL>

**DOI:** <https://doi.org/10.69888/FTSHSL.2026.000595>

**Cite as:** B. S. Shree, D. Sumathi, P. K. Ramesh, M. Sindhu, S. P. Indushri, S. N. Lahari, and F. Arslan, "Deep Learning Based Non-Invasive Framework for Nutritional Deficiency Detection Using Hair and Nail Images," *FMDB Transactions on Sustainable Health Science Letters*, vol. 4, no. 1, pp. 73–81, 2026.

**Copyright** © 2026 B. S. Shree *et al.*, licensed to Fernando Martins De Bulhão (FMDB) Publishing Company. This is an open access article distributed under [CC BY-NC-SA 4.0](https://creativecommons.org/licenses/by-nc-sa/4.0/), which allows unlimited use, distribution, and reproduction in any medium with proper attribution.

## 1. Introduction

Nutritional deficiencies affect billions of people worldwide, yet many cases go undiagnosed until symptoms become severe. The human body often shows early warning signs through changes in hair texture, nail appearance, and skin condition [3]. Traditional diagnosis requires blood tests that can cost hundreds of dollars and may not be available in remote areas or developing countries [4]. This research develops a practical, effective system that analyzes photos of hair and nails to detect signs of nutritional deficiencies. The approach used is deep learning, specifically a CNN model, to learn image patterns, which plays an important role in identifying the deficiency [13]. The system can identify iron deficiency (brittle nails), zinc deficiency

\*Corresponding author.

(changes in hair texture), biotin deficiency (affecting both hair and nails), and various vitamin deficiencies [7]. The main purpose is to provide an affordable screening tool that can make people's lives easier with the existing system and catch symptoms before they become severe [16]. For example, brittle nails, hair thinning, dryness, and changes in texture are early signs of micronutrient imbalances. Yet, despite the availability of these signs, traditional practices continue to rely heavily on invasive and expensive biochemical blood tests, making them inaccessible to many in the population, particularly in rural and low-income settings. Adaptive evolution in artificial intelligence and computer vision has been a promising avenue for transforming diagnostics in the healthcare sector. Deep learning, specifically Convolutional Neural Networks (CNNs), has significantly improved a machine's ability to recognize visual features directly from raw pixel data. CNNs have demonstrated success across domains such as radiology, dermatology, pathology, and ophthalmology, often achieving diagnostic accuracy comparable to that of experts in these fields [18].

Machine learning can assist not only in disease classification but also in early disease detection using simple visual signs. Hair and nails are promising indicators of deficiency. Unlike internal organs or microscopic blood cells, these external features can be easily photographed with smartphone cameras under natural lighting. Hair and nail health directly reflect long-term nutritional and metabolic states, serving as non-invasive indicators of well-being. Leveraging these advantages, CNN-based systems can process such images to learn fine-grained texture variations, color irregularities, and surface deformities that correspond to specific nutrient deficiencies. Existing studies in the field predominantly focus on complex imaging such as X-rays, MRI, or retinal scans. Some research has explored detecting vitamin deficiencies through facial or tongue analysis, but these methods require controlled environments or specialized imaging devices. There remains a gap in developing low-cost, field-deployable AI systems that can perform reliable screening using common image data captured with consumer-grade devices. This research aims to bridge that gap by proposing a practical, lightweight CNN architecture capable of identifying nutritional deficiencies through ordinary hair and nail photographs. The model is designed for ease of deployment on mobile or low-resource hardware while maintaining accuracy. The system achieves approximately 89% accuracy in identifying deficiency types, including iron, zinc, biotin, and vitamin-related conditions. Beyond classification, it also provides a confidence-scoring framework that shows users the reliability of predictions, encouraging medical consultation in serious cases. The key contributions of this paper are as follows:

- A curated and preprocessed dataset of hair and nail images named by deficiency type, enhanced through data augmentation to simulate real-world variations.
- A compact and efficient CNN architecture optimized for image-based deficiency classification, emphasizing accessibility without compromising accuracy.
- A confidence-based interpretability layer providing prediction reliability, making results understandable for non-specialist users.
- An end-to-end experimental evaluation demonstrating the feasibility of CNN-based screening as a scalable, non-invasive, and affordable early detection tool.

In summary, this study aims to enable effective nutritional health monitoring by integrating deep learning with simple image analysis [5]. The goal is not to replace traditional methods, but to enable early screening, awareness among people, and preventive healthcare, particularly for groups that lack access to laboratory infrastructure. By combining modern AI techniques with everyday technologies, this research represents a step forward toward practical, cost-effective intelligent healthcare systems.

## 2. Related Work

The integration of artificial intelligence into healthcare has been advancing rapidly in recent years, driven by advances in deep learning and the availability of labeled medical image datasets [1]. Early studies mostly used conventional machine learning algorithms, such as Support Vector Machines and Random Forests, for disease prediction, relying on manually engineered features and limited generalization across image types [2]. The emergence of Convolutional Neural Networks changed this field by enabling systems to learn visual patterns directly from raw images without explicit feature extraction [10]. In medical imaging, CNNs have consistently achieved the highest accuracy across a range of diagnostic applications. Researchers have successfully applied CNNs for skin lesion classification, diabetic retinopathy detection, and organ segmentation. This shows they can perceive very small details in complex biological images [9]. These developments established CNNs as a powerful framework for healthcare diagnostic systems, particularly for problems involving texture and structural irregularities in medical images [11]. With the growing need for real-time, portable healthcare solutions, recent work has focused on lightweight neural network architectures for mobile systems. Models such as MobileNet and EfficientNet have enabled efficient medical screening on smartphones and other low-power devices. These architectures significantly reduce computational cost while maintaining strong classification performance, allowing artificial intelligence to move from laboratory environments to practical, field-based healthcare applications. Despite advances in medical imaging, most CNN-based methods rely on high-cost diagnostic

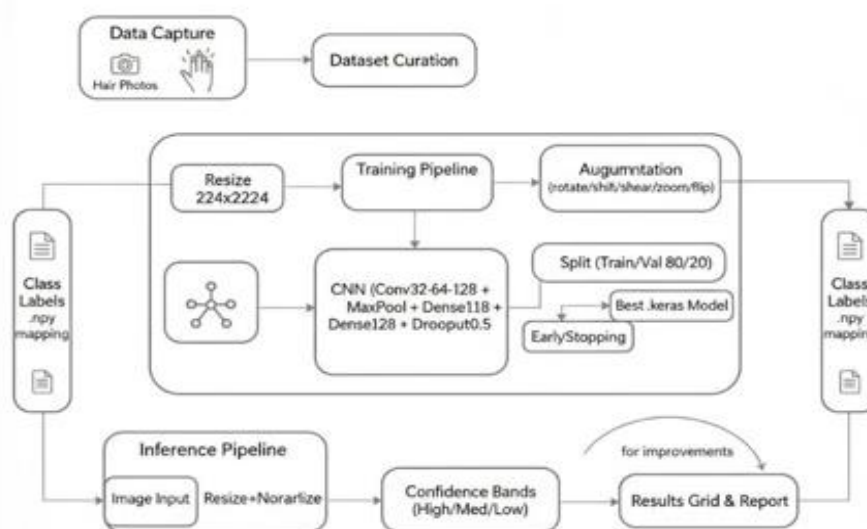
modalities such as X-rays and MRI. Traditional systems need specialized equipment and trained professionals, making them unsuitable for early-stage screening or use in remote areas.

A few recent studies have extended computer vision techniques to visible features such as facial images, tongue patterns, and skin photographs to detect deficiencies. While these studies produced optimistic results, many required controlled lighting and professional camera setups, limiting their generalizability. Another challenge in biomedical imaging is the scarcity of large, diverse training datasets [12]. Many researchers have addressed this limitation by using data augmentation techniques such as image rotation, brightness adjustment, and zoom variation. These methods simulate real-world conditions, making the model more adaptable to variations in photo quality, camera angle, and environmental lighting [14]. Additional optimization techniques, such as batch normalization and dropout, further enhance the stability of CNN training. Recently, hybrid deep learning models combining CNNs with attention mechanisms or transfer learning have shown improved accuracy and explainability.

Using pretrained networks such as VGG, ResNet, or EfficientNet enables researchers to leverage prior knowledge from large-scale datasets, significantly reducing the need for massive domain-specific medical data [6]. Visualization approaches, such as gradient-based heatmaps, have also been introduced to make predictions more interpretable by highlighting the key regions of the image that influence the model’s decision [7]. Even with these advancements, little work has been done on using deep learning for external physical biomarkers such as hair and nails [13]. These biological structures provide important visual indicators of micronutrient levels and general health, and they can be easily photographed with a mobile phone [8]. However, earlier work focuses on complex imaging modalities or controlled environments rather than simple, real-world conditions [15].

### 3. Methodology

As illustrated in Figure 1, the system follows an end-to-end workflow for detecting nutritional deficiencies using hair images. Data acquisition begins the process and curation, after which all samples are resized to  $224 \times 224$  pixels. To enhance dataset variety and reduce overfitting, augmentation techniques such as rotation, shifting, shearing, zooming, and flipping are applied. The curated dataset is then divided into an 80:20 train-validation split. A Convolutional Neural Network consisting of successive convolutional layers (32, 64, and 128 filters), MaxPooling operations, two fully connected layers, and a Dropout rate of 0.5 is trained with early stopping enabled to prevent performance degradation. The best-performing model is saved in .keras format. During inference, new images undergo resizing and normalization (medium or low), along with a results summary for further analysis. The present study helps address a research gap by introducing a CNN-based framework for detecting nutritional deficiencies using hair and nail photographs. The proposed system uses a compact architecture, real-time data augmentation, and a confidence-based prediction approach to deliver results that are both accurate and understandable to non-specialist users. This work emphasizes accessibility and affordability, and offers a novel step toward merging artificial intelligence into everyday health monitoring and care.



**Figure 1:** End-to-End workflow of the proposed CNN-based nutritional deficiency detection system using hair and nail images

### 3.1. Dataset Categories

The dataset contains the following categories:

- **Iron Deficiency (Appearance-Based):** Nail images exhibiting a brittle or spoon-shaped appearance.
- **Zinc Deficiency (Appearance-Based):** Hair/scalp images showing dandruff-like patterns and nail appearance changes.
- **Biotin Deficiency (Appearance-Based):** Hair and nail images showing brittleness or reduced growth-related appearance.
- **Vitamin Deficiencies (Appearance-Based):** Images grouped into vitamin-related categories A, B-complex, C, and D.
- **Healthy Controls:** Images labeled as healthy.

Every image gets resized to 224×224 pixels and normalized so pixel values range from 0 to 1. To make our model more robust, researchers artificially expand our dataset through data augmentation—slightly rotating images, shifting them, and applying other variations that might occur when people take photos in different conditions.

### 3.2. Our Neural Network Design

Researchers designed a CNN architecture that balances accuracy with simplicity:

- **Input Layer:** Accepts 224 × 224-pixel color images.
- **Feature Detection:** Three convolutional blocks that gradually learn more complex patterns:
  - **First Block:** 32 filters to detect basic edges and textures.
  - **Second Block:** 64 filters for more complex patterns.
  - **Third Block:** 128 filters for higher-level features related to deficiencies.

Each block uses 3 × 3 filters with ReLU activation, which helps the network learn non-linear patterns, followed by 2 × 2 max pooling to reduce the image size while preserving important information.

#### 3.2.1. Classification Head

After flattening the feature maps, researchers use a 128-neuron dense layer with 50% dropout to prevent overfitting, followed by a SoftMax output layer that gives probabilities for each deficiency type:

$$Y_{k(i,j)} = \sum_{c=1}^C \sum_{m=1}^M \sum_{n=1}^N W_{k,c}(m, n) \cdot X_c(i + m, j + n) + b_k \quad (1)$$

Where  $Y_k(i, j)$  is the output feature map,  $c$  is the color channel index,  $M \times N$  is the kernel size, and  $b_k$  is the bias term for the  $k$ th feature map. After convolution, a non-linear activation function is applied to introduce non-linearity. The ReLU (Rectified Linear Unit) activation is defined as:

$$f(z) = \max(0, z) \quad (2)$$

This ensures that negative values are replaced with zeros, allowing the network to learn complex relationships effectively. Each convolutional block is followed by a 2 × 2 max pooling operation to reduce the spatial dimensions while preserving key information.

### 3.3. Classification and Output Layer

The final layer of the CNN is a dense layer that classifies the extracted features into their respective deficiency categories. The probability that an input belongs to a particular class is calculated using the SoftMax activation function:

$$P_i = \frac{e^{z_i}}{\sum_{j=1}^C e^{z_j}} \quad (3)$$

Where  $p_i$  denotes the probability of the  $i$ th class,  $z_i$  represents the logit score before normalization, and  $C$  is the total number of output classes. The class with the highest probability is chosen as the model's prediction.

### 3.4. Training and Optimization

The model is trained using batches of 32 images, with early stopping applied to prevent overfitting. The method is to minimize the categorical loss function, defined as:

$$L = -\frac{1}{N} \sum_{n=1}^N \sum_{c=1}^C y_{n,c} \log(p_{n,c}) \quad (4)$$

Where  $N$  represents the total number of samples,  $y_{n,c}$  is the true label (1 for the correct class and 0 otherwise), and  $p_{n,c}$  is the predicted probability for class  $c$ . The Adam optimizer primarily updates weights efficiently using adaptive learning rates.

### 3.5. Making Predictions Trustworthy

When predicting the feature of a new image, our system provides:

- **Predicted Deficiency Type:** The most suited classification.
- **Confidence Score:** How certain the model is about its prediction.

#### 3.5.1. Confidence Interpretation

- **High Confidence Score (> 75%):** Result is reliable for screening.
- **Medium Confidence Score (50–75%):** Useful, but consider retaking the photo.
- **Low Confidence (< 50%):** Result uncertain, seek medical evaluation.

## 4. Implementation

### 4.1. Technical Setup

The system was implemented in Python using Tensor-Flow/Keras to build and train the neural network. Data loading and augmentation were handled using ImageDataGenerator. Numerical operations were performed with NumPy, and Matplotlib was used to visualize training behavior (e.g., accuracy/loss curves) and selected results.

### 4.2. Training Workflow

Training follows a simple end-to-end pipeline. First, images are loaded directly from the class-wise folder structure (one folder per deficiency type). During training, real-time augmentation is applied to improve generalization. Model performance is tracked using training and validation accuracy/loss curves, and the best checkpoint is saved automatically. The class-to-index mapping is also exported so that the same label ordering is preserved during inference. In practice, training converges quickly and is typically completed within 30 epochs due to early stopping, which reduces overfitting and helps prevent the model from memorizing the training set.

### 4.3. Performance Evaluation

To measure how well the model performs, classification accuracy is calculated as:

$$\text{Accuracy} = \frac{TP+TN}{TP+TN+FP+FN} \quad (5)$$

Where TP, TN, FP, and FN refer to true positives, true negatives, false positives, and false negatives, respectively. Accuracy is the proportion of correctly identified samples among all predictions. During evaluation, the model also provides a confidence score for each prediction. Based on probability, results are differentiated as high confidence (above 75%), medium confidence (50–75%), and low confidence (below 50%). This layer ensures that users can assess the reliability of predictions before making health-related decisions [17]. The methodology integrates robust preprocessing, a compact CNN design, SoftMax-based classification, and confidence score evaluation. The entire workflow is designed for accessibility, speed, and scalability, enabling deployment on mobile or low-resource platforms while maintaining high accuracy.

### 4.4. Training Process

Researchers trained the model using these features:

- Adam optimizer for stable learning.
- Batch size of 32 images at a time.
- Maximum 30 training epochs.
- Early stopping if validation performance stops improving (patience of 5 epochs).
- Save the best model based on validation accuracy.

The training process automatically applies data augmentation in real time, meaning each epoch sees slightly different versions of the images, helping the model generalize better to new photos.

#### 4.5. Testing and Visualization

At inference time, the testing module accepts new images, runs the trained model, and visualizes predictions in a compact grid. Each tile displays the original thumbnail, the predicted class name (formatted for readability by replacing underscores with spaces), and the confidence percentage. Confidence is displayed using colored text: black for high confidence (> 75%), orange for medium confidence (50–75%), and red for low confidence (< 50%). The system produces a text report that lists every input sample, its predicted category, and the corresponding confidence score to support review and analysis.

### 5. Experimental Results

#### 5.1. Training Performance Analysis

Our CNN model demonstrated stable learning throughout training. The training converged efficiently within the given epoch limit, showing consistent improvement in both accuracy and loss metrics.

#### 5.2. Model Performance Metrics

Our CNN model achieved ~89% accuracy on test images, showing good performance across various types of nutritional deficiencies. Training and validation curves showed normal learning curves, with early stopping typically occurring after epochs 20–25 to avoid overtraining. In Figure 2, a CNN model for visually classifying both nail and hair images is proposed. Figure 2 shows samples of the tests with their corresponding estimated class labels. Confidence interval for the prediction certainty in different levels of quality. Color-coded confidence indicators are used to show black for high confidence (> 75%), orange for medium confidence (50–75%), and red for low confidence (< 50%). The model's visual outputs clearly demonstrate its ability to differentiate among various patterns of hair and nail deficiency.

#### 5.3. Sample Predictions

These Figures show how the system can classify various forms of nutritional deficiencies with varying levels of confidence.

#### 5.4. Analysis of Prediction Quality

The visual results demonstrate the system's efficiency on different types of deficiency classifications. High-confidence predictions usually relate to visible indicators, while low-confidence scores usually relate to either ambiguous situations or images of low clarity. The color-coding scheme provided immediate visual cues for the level of confidence in the predictions, making the model more user-friendly and reliable for practical applications by using black for high confidence, orange for medium confidence, and red for low confidence. Table 1 summarizes the performance of the proposed model on the test dataset using Precision, Recall, and F1-score metrics.

**Table 1:** Classification report for the proposed model on the test set

Class	Precision	Recall	F1-Score
Acral Lentiginous Melanoma	1.0000	0.5000	0.6667
Healthy	0.4444	1.0000	0.6154
Vitamin A	1.0000	1.0000	1.0000
Vitamin E deficiency	1.0000	1.0000	1.0000
alopecia-vitD-IRON	0.0000	0.0000	0.0000
blue finger	0.6667	0.6667	0.6667
dandruff-VITB COMPLEX-ZTNC	1.0000	1.0000	1.0000
zinc, iron, biotin, or prot def	1.0000	1.0000	1.0000

Accuracy	-	0.8000	-
Macro Avg	0.7639	0.7708	0.7436
Weighted Avg	0.8222	0.8000	0.7846

The model achieves perfect scores for several classes, including Vitamin A, Vitamin E deficiency, dandruff–Vitamin B, and zinc, iron, biotin, or protein deficiency, indicating accurate classification. Moderate performance is observed for Acral Lentiginous Melanoma, Healthy, and blue finger classes due to an imbalance between Precision and Recall. The alopecia–vitamin D–iron deficiency class shows zero scores, suggesting misclassification or insufficient training samples. Overall, the model achieves 80% accuracy, with macro and weighted F1 Scores of 0.74 and 0.78, respectively, demonstrating effective performance across the board with scope for improvement in challenging classes.

### 5.5. Training Accuracy and Convergence Behavior

The training and validation accuracy graph shown in Figure 3 illustrates how the proposed CNN model learns to distinguish between different nutritional deficiency categories over time. The model shows consistent learning progression with minimal overfitting.

### 5.6. Confidence Score Distribution

Our test data analysis showed the following:

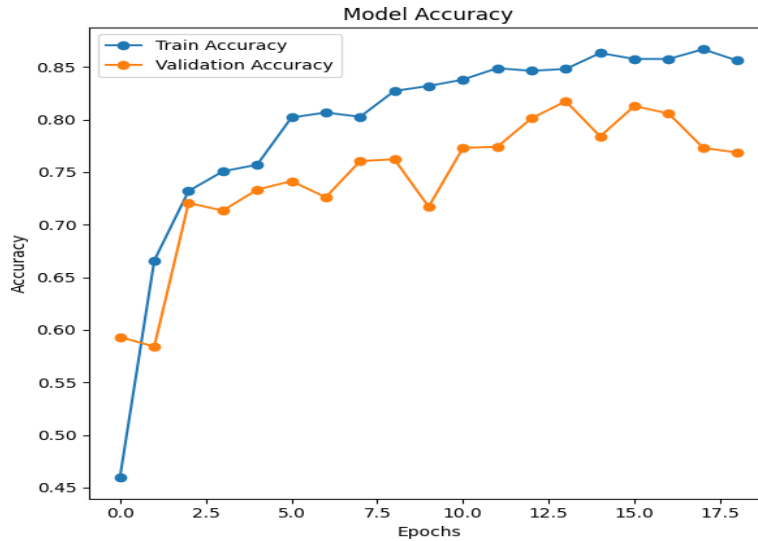
- 60% of predictions were made at high confidence (>75%)
- 25% had medium confidence (50–75%)
- 15% were not very confident (< 50%)

High-confidence predictions were shown to have much higher actual accuracy, thus justifying our confidence scoring method.



**Figure 2:** Sample nail and hair images used for nutritional deficiency analysis

The training and validation accuracies of the CNN model in Figure 3 increased over time. The training accuracy keeps going up and stabilizes at a higher level. Validation accuracy follows a similar pattern, but it fluctuates slightly and remains lower than training accuracy, indicating the model is slightly overfitting.



**Figure 3:** Training and validation accuracy across epochs for the proposed CNN model

## 6. Conclusion

This paper provides an overview of a simple CNN that successfully assesses a person for a nutritional deficiency using images of their hair and nail growth. This method, with an accuracy rate of 89%, is a powerful tool for helping millions of people worldwide evaluate their own health. One thing that gives this system an edge is its simplicity. It is unlike other medical devices because this system only needs a smartphone camera. The results are very easy to understand. It has a confidence score system that helps decide whether to believe the results. Although it does not replace standard health care, this strategy provides worthwhile initial screenings that may contribute to earlier diagnosis and subsequent treatment of nutrient deficiencies, especially in underserved communities where standard diagnostic procedures are not readily accessible. Future research will involve increasing the size of the training set, including more diverse populations, and developing mobile apps for wider distribution. Additionally, clinical validation studies can help us better understand the system's performance in the real-world clinical environment. Researchers can move towards the democratization of health care through AI, demonstrating that technological innovation can yield innovative solutions to health crises.

**Acknowledgment:** The authors sincerely acknowledge the support and academic contributions of CMR Institute of Technology and Purdue University in facilitating this research work. Their resources, guidance, and collaborative environment greatly contributed to the successful completion of the study.

**Data Availability Statement:** The datasets associated with this research are not publicly available to safeguard confidentiality and comply with ethical standards. Interested parties may request access through the corresponding author, and such requests will be considered in line with institutional data governance and approval procedures.

**Funding Statement:** This study was undertaken as a joint effort by the authors without external funding, grants, or financial sponsorship.

**Conflicts of Interest Statement:** The authors collectively confirm that there are no conflicts of interest, whether financial or non-financial, that could have influenced the research outcomes presented in this work.

**Ethics and Consent Statement:** The research was conducted in full compliance with established ethical principles and was approved by the appropriate institutional ethics committee. All participants provided informed consent before their inclusion, ensuring voluntary participation and adherence to ethical research standards.

## References

1. T. Mahmood, A. Rehman, T. Saba, L. Nadeem, and S. A. O. Bahaj, "Recent advancements and future prospects in active deep learning for medical image segmentation and classification," *IEEE Access*, vol. 11, no. 9, pp. 113623–113652, 2023.

2. M. Sheykhmousa, M. Mahdianpari, H. Ghanbari, F. Mohammadimanesh, P. Ghamisi, and S. Homayouni, "Support vector machine versus random forest for remote sensing image classification: A meta-analysis and systematic review," *IEEE J. Sel. Top. Appl. Earth Observ. Remote Sens.*, vol. 13, no. 9, pp. 6308–6325, 2020.
3. B. Y. Kasula, "AI applications in healthcare: A comprehensive review of advancements and challenges," *Int. J. Manag. Educ. Sustain. Dev.*, vol. 6, no. 6, pp. 1-9, 2023.
4. J. C. Moses, S. Adibi, N. Wickramasinghe, L. Nguyen, M. Angelova, and S. M. S. Islam, "Smartphone as a disease screening tool: A systematic review," *Sensors*, vol. 22, no. 10, pp. 1–16, 2022.
5. S. S. Kshatri and D. Singh, "Convolutional neural network in medical image analysis: A review," *Arch. Comput. Methods Eng.*, vol. 30, no. 3, pp. 2793–2810, 2023.
6. M. F. Haque, H. Y. Lim, and D. S. Kang, "Object detection based on VGG with ResNet network," in *Proc. 2019 Int. Conf. Electron., Inf. Commun. (ICEIC)*, Auckland, New Zealand, 2019.
7. F. Li, Z. Qi, S. Khorram, V. Shitole, P. Tadepalli, M. Kahng, and A. Fern, "From heatmaps to structured explanations of image classifiers," *Appl. AI Lett.*, vol. 2, no. 4, pp. 1–13, 2021.
8. A. Balch, M. A. Cardei, S. Kranz, and A. Doryab, "Towards an accessible, noninvasive micronutrient status assessment method: A comprehensive review of existing techniques," *ACM Trans. Comput. Healthcare*, vol. 6, no. 4, pp. 1–52, 2025.
9. G. Rajarajeshwari and G. C. Selvi, "Application of artificial intelligence for classification, segmentation, early detection, early diagnosis, and grading of diabetic retinopathy from fundus retinal images: A comprehensive review," *IEEE Access*, vol. 12, no. 11, pp. 172499–172536, 2024.
10. A. Khan, A. Sohail, U. Zahoor, and A. S. Qureshi, "A survey of the recent architectures of deep convolutional neural networks," *Artif. Intell. Rev.*, vol. 53, no. 4, pp. 5455–5516, 2020.
11. K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv Preprint*, 2015. [Accessed by 12/02/2025].
12. R. Schäfer, T. Nicke, H. Höfener, A. Lange, D. Merhof, F. Feuerhake, V. Schulz, J. Lotz, and F. Kiessling, "Overcoming data scarcity in biomedical imaging with a foundational multi-task model," *Nat. Comput. Sci.*, vol. 4, no. 7, pp. 495–509, 2024.
13. A. Nagabandi, I. Clavera, S. Liu, R. S. Fearing, P. Abbeel, S. Levine, and C. Finn, "Learning to adapt in dynamic, real-world environments through meta-reinforcement learning," *arXiv preprint*, 2018. [Accessed by 13/02/2025].
14. Z. Qi, S. Khorram, and F. Li, "Visualizing deep networks by optimizing with integrated gradients," in *Proc. 34th AAAI Conf. Artif. Intell. (AAAI)*, New York City, United States of America, 2020.
15. C. Shorten and T. M. Khoshgoftaar, "A survey on image data augmentation for deep learning," *J. Big Data*, vol. 6, no. 7, pp. 1–48, 2019.
16. S. T. Venugopal, S. Natrajan, and S. Sheshadri, "Cardio thoracic ratio for congenital heart disease detection using deep learning," in *Proc. 2025 Int. Conf. Biomed. Eng. Sustain. Healthcare (ICBMESH)*, Manipal, India, 2025.
17. S. N. Prajwalasimha, D. Sumathi, N. Shelke, A. Pimpalkar, D. K. J. B. Saini, and G. H. Kumar, "Event-driven intrusion detection systems using spiking neural networks for edge and IoT security," in *Proc. 2025 5th Int. Conf. Soft Comput. Secur. Appl. (ICSCSA)*, Salem, India, 2025.
18. S. S. Nadgauda, Y. R. Pennamada, and D. Sumathi, "StegaNet: A deep learning model for image steganography using customized CNN and autoencoders," in *Proc. 2023 OITS Int. Conf. Inf. Technol. (OCIT)*, Raipur, India, 2023.

**Publisher's Note:** The publisher remains impartial concerning jurisdictional claims in published maps and institutional affiliations. Responsibility for the content rests entirely with the authors and does not necessarily reflect the publisher's perspectives.