

Non - Invasive Detection of Vitamin Deficiency Using Resnet and Image Processing

P. Srija Vagdevi¹, T.Sai Sampath², G.Sai Madhuri³, V.Sachin⁴

^{1,2,3,4}Students, Department of Computer Science and Engineering, Lendi Institute of Engineering and Technology (Autonomous), Andhra Pradesh, India

A.V.D. N Murthy, M.Tech, Assistant Professor, Department of Computer Science and Engineering, Lendi Institute of Engineering and Technology (Autonomous), Andhra Pradesh, India

paidasrija@gmail.com¹, saisampath672@gmail.com², govardhanagirisaimadhuri717@gmail.com³, sachinvudiga14082003@gmail.com⁴

ABSTRACT - Vitamin Deficiencies are not identified and addressed in a timely manner, can cause major health problems. Blood tests, which can be intrusive, time-consuming, and resource-intensive, are frequently used in traditional techniques of diagnosing vitamin deficiencies. In our work, we suggest a non-invasive method for detecting possible vitamin deficiencies from facial photos or other biomarkers such changes in the texture of the skin and nails by utilizing deep learning and image processing techniques. Utilizing pre-trained models and convolutional neural networks (CNNs) such as RESNET50, we extract important features from these images and link them with recognized signs of inadequacies. The system is trained using a collection of tagged photos of healthy controls and people with different vitamin deficits. Using thorough picture analysis, the model detects visual indicators of inadequacies such as skin discoloration, pallor, anomalies of the eyes, and changes in hair. We are predicting deficiencies for vitamins A, B2, B3, B5, B6, B7, B12, C, D, E, and K. However, we are not detecting deficiencies for vitamins B1 and B9 because the textures we have used to extract features are not suitable for identifying these particular vitamins. Our initial findings indicate that we can anticipate deficiency levels with a promising degree of accuracy. This could provide a rapid, easily accessible, and non-invasive early screening tool that could support preventative healthcare initiatives and standard diagnostic techniques.

Key Words: : Convolutional Neural Networks (CNNs), ResNet50, Deep Learning, Image Processing, Feature Extraction

1. INTRODUCTION

Vitamin deficiencies significantly impact health, causing symptoms ranging from fatigue to cognitive problems. Early detection is crucial, but traditional blood tests are invasive, expensive, and often delay treatment. Our research proposes a non-invasive alternative using deep learning and image processing to identify vitamin deficiencies through visual markers in skin, nails, and facial features.

We implement ResNet50 (Residual Network) architecture, chosen for its exceptional image recognition capabilities through skip connections that efficiently learn complex patterns. The system analyses specific visual indicators of vitamin deficiencies, including nail ridges and skin changes.

To strengthen the model's training, we've developed a comprehensive dataset featuring vitamin deficiency symptoms. We employ image preprocessing techniques such as contrast adjustment, edge enhancement, and normalization to improve feature visibility and overall performance.

The system undergoes rigorous evaluation using accuracy, precision, recall, and F1-score metrics. Our goal is to deliver an accessible, user-friendly solution for early vitamin deficiency detection in healthcare settings and for individuals seeking preventive care.

2. LITERATURE SURVEY

Deep Learning for Tomographic Image Reconstruction: The use of convolutional neural networks (CNNs) and other deep learning techniques in medical imaging is covered in this article. It emphasizes how these strategies deal with issues that are relevant to

non-invasive diagnostic procedures, like noise and data sparsity [1].

Medical Image Computing: This source offers a summary of segmentation methods used in medical imaging, such as CNNs and ResNet models. It highlights the use of these models in medical image analysis, which is essential for using image processing to identify diseases like vitamin deficiency [2].

General Overview of Deep Learning in Medical Imaging: This paper provides a thorough overview of deep learning applications in medical imaging, addressing different architectures and their efficacy in various diagnostic tasks. It is a fundamental resource for comprehending the integration of deep learning in medical diagnostics [3].

Use of Deep Learning in Nutritional Deficiency Detection: This work investigates the use of facial feature analysis and deep learning models to identify nutritional deficiencies. It illustrates how non-invasive techniques can be used to use image analysis to detect vitamin deficiencies [4].

ResNet's Effect on Medical Image Analysis This study explores the use of Residual Networks (ResNet) in medical picture processing, emphasizing how well they can identify intricate patterns and raise diagnostic precision. It emphasizes how useful ResNet is for initiatives that use imaging to identify medical disorders [5].

3.DATASET COLLECTION

As we know that collection of data is very important for starting our project ,we started data collection from scratch.to ensure that we had a complete dataset for our project, our team collected a variety of images from open accessible sources, such as Google. We used data augmentation techniques to improve the dataset's quality and variability. Transformations like rotation, flipping, scaling, cropping, brightness modifications, and noise addition were among these methods. We sought to increase the model's robustness and generalization capabilities across various circumstances by expanding the dataset, which would lower the likelihood of overfitting.We examined 19,392 images of vitamin deficiency symptoms with total of 13 class labels applied. These images are divided into 13,480 training images, 1978 testing images and 3934 validation images. Given only the picture of the vitamin deficiency, we try

to predict the vitamin deficiencies that corresponds to each class label.

The Table-1 shows all 13 classes and the number of images in each class.

CLASS LABELS	TRAIN	TEST	VALID
Vitamin A	795	115	227
Vitamin B2	1890	271	540
Vitamin B3	444	64	127
Vitamin B5	241	35	69
Vitamin B6	1131	162	323
Vitamin B7	833	119	238
Vitamin B12	1422	204	406
Vitamin C	302	44	86
Vitamin D	2089	299	597
Vitamin E	1523	219	435
Vitamin K	788	114	225
Healthy	1101	145	289
Wound	1011	187	372

4.METHODS AND TECHINQUES:


4.1ResNet50:

A deep learning architecture called ResNet (Residual Network) was created to address the challenge of training extremely deep neural networks. The vanishing gradient problem, which occurs when gradients are too small during backpropagation and render training useless, frequently affects deeper neural networks. In order to solve this problem, ResNet uses skip connections, also known as residual connections, which let the network transfer data straight to later layers while avoiding some tiers.

The model can train deeper networks without experiencing performance deterioration because to these skip connections, which also help to preserve crucial properties. ResNet makes it possible to train models with hundreds or even thousands of layers, which greatly increases accuracy in challenging image identification applications. In computer vision applications like object identification, picture classification, and medical imaging, ResNet has been widely employed due to its capacity to learn deep and meaningful features. The most widely used variants are ResNet50, ResNet101, and ResNet152, the number represents the total number of layers in the model.


Essential Elements:

1. Residual blocks: ResNet is constructed using these blocks. Every block has two paths: the highway, which is the primary road, and the on-ramp, which is the skip path. Whereas the skip path keeps the prior data, the main path picks up new features.



2. Skip Connections: With these connections, the model can learn efficiently even if it skips some levels.

3. Batch normalization: Is a technique that enhances the model's performance and helps stabilize the learning process.



ResNet50 Model Architecture

4.2 ADAM:

The Adam Optimizer is a well-known deep learning optimization technique that assists in modifying a model's parameters in order to reduce the loss function. It is an expansion of the stochastic gradient descent (SGD) technique, which allows the model to converge more quickly and precisely by adjusting the learning rate for each parameter according to the gradient's magnitude. Adam Optimizer normalizes gradients to stabilize the learning process and adds momentum to assist the model escape local minima and converge more quickly. Deep learning models can enhance overall stability, converge more quickly, and require less hyperparameter adjustment when Adam Optimizer is used. Step size, momentum, adaptive learning rate, and

numerical stability are all regulated by the algorithm's learning rate, beta1, beta2, and epsilon, among other important hyperparameters.

4.3 Global Average Pooling (GAP):

Convolutional Neural Networks (CNNs), Global Average Pooling 2D (GAP 2D) is a pooling operation that reduces the spatial dimensions (height and breadth) of feature maps while maintaining significant features. One value per feature map is produced by averaging the values of all the feature maps throughout their whole geographical grid. By drastically lowering the network's parameter count, this procedure improves the model's computational efficiency and helps avoid overfitting. As the model grows increasingly invariant to object positions within the image, GAP 2D enhance In generalization by emphasizing global properties rather than exact spatial locations.

The network can still comprehend the general patterns in the image thanks to GAP 2D's retention of important global information, which is particularly helpful for tasks like image categorization. GAP 2D, in short, provides a straightforward method of compressing feature maps while preserving important information, resulting in a more effective and broadly applicable model.

4.4 Dropout:

The Dropout strategy in ResNet randomly removes neurons during training in order to avoid overfitting. By randomly setting a subset of neurons' output to zero, Dropout, which is applied after convolutional and batch normalization layers but before the ReLU activation function, pushes the network to learn different representations of the data. This results in less overfitting, better generalization, and greater robustness by lowering the network's reliance on any one neuron. ResNet typically employs a dropout rate of 0.5, which indicates that 50% of neurons are removed during training.

4.5 Data Augmentation:

Data augmentation is the process of introducing random alterations to the existing data in order to artificially expand the size of a training dataset. Images can be flipped, rotated, cropped, and their brightness and contrast changed. Audio files can also be distorted or noise added. This exposes the model to a greater range

of data, which enhances its performance and generalization to new, unknown data by helping it learn to identify patterns and features more successfully. This helps to avoid overfitting and makes the model more resilient and flexible, which is particularly helpful when working with tiny datasets.

Furthermore, by giving the model a wider variety of instances, data augmentation can lessen bias in the model and potentially enhance its capacity to manage noise and actual variances.

4.6 Callback Techniques:

The callback techniques are used in our project to optimize the training process, improve model performance, and make the most efficient use of computational resources. Here's how each one directly supports our vitamin deficiency detection model. Types:

1. Model Checkpoint Callback
2. Early Stopping Callback
3. ReduceLROnPlateau Callback

4.6.1 Model Checkpoint Callback

When the model performs best (i.e., has the lowest validation loss), it is saved. This guarantees that, in our project, you will keep the model that performs the best for detecting vitamin deficiencies, even if subsequent epochs result in overfitting or worse performance. It permits you to reload the optimal model for assessment, deployment, or fine-tuning without losing progress.

4.6.2 Early Stopping Callback

After a predetermined number of epochs, the system ends training when no improvement is observed, continuously monitoring metrics (such as validation loss). This helps prevent overfitting, particularly when working with synthetic image data that might not fully represent real-world variance. Additionally, by terminating training early when the model has already reached optimal performance, it conserves computational resources and reduces development time.

4.6.3 ReduceLROnPlateau Callback

When a metric stops improving, it lowers the learning rate so the model can update weights more precisely. This is crucial for our project since the model can plateau while being trained. A superior local optimum may be

overlooked by a high learning rate. It aids in improving the model's convergence, which is crucial for optimizing a pretrained model.

4.7 Activation Functions:

4.7.1 ReLu

One of the most popular activation functions in deep learning, particularly in convolutional and thick layers, is the ReLu (Rectified Linear Unit) function. By simply producing the input directly if it is positive and zero otherwise, it adds non-linearity to the model; formally, this is expressed as $\text{ReLu}(x) = \max(0, x)$. ReLu is well-liked due to its computational efficiency and ability to alleviate the vanishing gradient issue, which facilitates quicker training and improved deep neural network performance.


4.7.2 SOFTMAX

For multi-class classification tasks, the output layer of deep learning models frequently uses the SoftMax activation function. It creates a probability distribution from a vector of raw prediction scores (logits), where all values add up to one and each value ranges from 0 to 1. This makes it easier to understand the model's output, which is the likelihood of each class. By normalizing the output and exponentiating each logit, it mathematically highlights the greatest values while suppressing the lower ones. SoftMax is frequently used in conjunction with categorical cross-entropy loss to efficiently train classification models.

5. PROPOSED METHODOLOGY:


Our system is now designed to predict a group of vitamin deficiencies rather than just one specific deficiency. Unlike previous methodologies, which focused on detecting only one or two particular vitamin deficiencies, our model is capable of predicting deficiencies in multiple vitamins, including A, B2, B3, B5, B6, B7, B12, C, D, E, and K. The proposed algorithm begins by collecting a dataset of facial photos and other biomarkers (skin texture, nails) from These images are preprocessed, resized, and normalized for input. A pre-trained Convolutional Neural Network (CNN) like ResNet50 is used to extract key features from these images. The extracted features are then used to train a classifier model to predict vitamin deficiencies based on visual indicators like skin discoloration and hair changes. The model is evaluated using accuracy and other metrics, and once

trained, it can predict deficiencies in new images. This provides a non-invasive, accessible, and efficient screening tool for early detection of vitamin deficiencies.





Model Architecture

6. BLOCK DIAGRAM:



7.EXPERIMENTAL RESULTS & DISCUSSIONS:

The vitamin deficiency detection model, based on ResNet-50, showed steady improvement during training. Both the training and validation accuracies gradually increased, while the training and validation losses consistently decreased over epochs. Although the learning trend indicates that the model is effectively extracting features and learning classification patterns. The model successfully differentiates between multiple vitamin deficiency categories, and with further training, more data, and fine-tuning, it has strong potential for highly accurate prediction of vitamin deficiencies from facial or nail images.



The both graphs shows the model accuracy and model loss of our model.

FACTORS USED TO CALCULATE ACCURACY:

1.ACCURACY

Accuracy measures the overall correctness of the model across all classes present in our project. It tells what percentage of total predictions were correct.

$$\text{Accuracy} = \frac{TP+TN}{TP+TN+FP+FN}$$

2.PRECISION

Precision measures how many of the positively predicted cases were actually correct. It is important when false positives need to be minimized.

$$\text{Precision} = \frac{TP}{TP+FP}$$

3.RECALL

Recall measures how many of the actual positive cases were correctly identified by the model. It is important when missing true cases (false negatives) is risky.

$$\text{Recall} = \frac{TP}{TP+FN}$$

4. F1-SCORE

F1-score balances precision and recall using their harmonic mean. It is useful when you need a balance between correctly predicted and all actual positives.

$$F1Score = 2 \times \text{Precision} \times \text{Recall} / (\text{Precision} + \text{Recall})$$

Metric	Formula	Meaning
Accuracy	$(TP + TN) / (TP + TN + FP + FN)$	Overall correctness
Precision	$TP / (TP + FP)$	Correctness of positive guesses
Recall	$TP / (TP + FN)$	Finding all actual positives
F1-Score	$2 \times (\text{Precision} \times \text{Recall}) / (\text{Precision} + \text{Recall})$	Balance of Precision and Recall

Factor values of our model:

METRICS	VALUES OBTAINED
Accuracy	0.9216
Precision	0.9217
Recall	0.9216
F1-Score	0.9216

8.CONCLUSION:

In conclusion, this project highlights the potential of deep learning and image processing in revolutionizing healthcare diagnostics by giving 90-95% By providing a non-invasive, efficient method to detect vitamin deficiencies through visual biomarkers, it addresses the limitations of traditional blood tests. This innovative solution can aid in early diagnosis and timely intervention. It paves the way for more accessible and user-friendly health monitoring tools in the future.

Future Scope: Our system is designed to find vitamin deficiency for the multiple deficiencies like vitamin A,B2,B3,B5,B6,B7,B12,C,D,E,K.by using the textures like **(eyes,nails,tongue,hair,lips,skin)** but we are not detecting deficiencies for vitamins B1 and B9 because the images relevant to these vitamin deficiencies do not come under our proposed metrics.

9.REFERENCES:

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep Residual Learning for Image Recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 770–778. <https://doi.org/10.1109/CVPR.2016.90>[1]

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet Classification with Deep Convolutional Neural Networks. In Advances in Neural Information Processing Systems (NeurIPS), 25, 1097–1105.[2]

World Health Organization (WHO). (2023). Micronutrient deficiencies. Retrieved from: <https://www.who.int/health-topics/micronutrients>[3]

Prateek, A., & Jain, R. (2020). Non-Invasive Prediction of Vitamin Deficiency Using Image Processing and Deep Learning. International Journal of Computer Applications, 176(12), 1–6.[4]

Chollet, F. (2017). Xception: Deep Learning with Depthwise Separable Convolutions. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 1251–1258.P[5]

OpenCV. (2024). Image Processing Library. Retrieved from: <https://opencv.org/>[6]

TensorFlow. (2024). TensorFlow: An end-to-end open source machine learning platform. Retrieved from: <https://www.tensorflow.org/>[7]

Keras Documentation. (2024). Applications - ResNet50. Retrieved from: <https://keras.io/api/applications/resnet/#resnet50-function>[8]

NIH (National Institutes of Health). (2022). Vitamin and Mineral Supplement Fact Sheets. Retrieved from: <https://ods.od.nih.gov/factsheets/list-all/>[9]

P. Sharma et al. (2021). AI-Based Skin and Nail Analysis for Detecting Nutritional Deficiencies: A Clinical Approach. Journal of Biomedical Informatics, 117, 103771.[10]