IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)
An International Open Access, Peer-reviewed, Refereed Journal

A RESNET-BASED FRAMEWORK FOR VITAMIN DEFICIENCY DETECTION FROM VISIBLE SYMPTOMS

[1]Geetha Sri Yendrapati, [2]Yasasvi Sai Sree Toleti, [3]Nehapriya Yarlagadda,[4]Mani Kumar Chari Sriramadasu, [5]Sanjay Saikam

[1]Assistant Professor, [2]Student, [3]Student, [4]Student, [5]Student
[1]Department of Computer Science and Engineering,
[1]Seshadri Rao Gudlavalleru Engineering College, Gudlavalleru, India

Abstract: Vitamin deficiencies represent one of the most common health problems in many countries of the world, especially where access to medical test facilities is poor. These deficiencies could take the form of impaired immunity, a sight problem, skin diseases, anemia, or poor physical development. Conventional techniques for its diagnosis primarily rely on blood tests, which are very accurate but sometimes invasive, expensive, and ill-suited for mass screening or screening in remote areas. This demands the development of other methods of diagnosis that are non-invasive in nature and yet would facilitate early detection. This project is aimed at the development of an AI-driven system for the detection of vitamin deficiencies using images with visible symptoms from eyes, nails, skin, tongue, and hair by deep learning. In this approach, transfer learning applies the ResNet152V2 model for feature extraction, which uses the attention mechanism to help the model focus on most of the significant parts of the image that are associated with deficiency-related visual patterns, improving the detection of very small manifestations of symptoms. A deep learning approach using a custom dataset of images of different body areas and corresponding labels is used to train and test the proposed system. The proposed model will be able to detect multiple vitamin deficiencies and give accurate predictions. The performance of the proposed model tested using common evaluation metrics such as accuracy(91.4%), precision(91%), recall(90.7%), and F1-score(90.8%), which will show the efficiency of the proposed classification model. The proposed model is expected to be used as a supporting diagnostic tool and will work to decrease the dependence on laboratory work. The proposed model will also be non-invasive and will be easily deployable on a large scale.

Index Terms - Vitamin Deficiency Detection, Deep Learning, ResNet152V2, Transfer Learning, Image Classification, Visible Symptom Analysis, Remote Healthcare.

I. Introduction

Vitamin deficiencies are a very widespread yet serious concern in public health and fall under the category of nutritional disorders that affect human health adversely in a significant manner. Vitamins are micronutrients that play a vital part in enabling humans to possess normal metabolic processes and also in developing immunity levels in a human being along with contributing to proper mental functions and development. As a result of a lack of proper availability of these micronutrients to vital bodily functions in a human being, a host of health-related problems arise in a person that includes anemia, vision impairments, skin-related ailments, poor immunity levels in a human being, mental fatigue in a person's mind, delayed physical development in a child's body, and in worse cases, a person may also undergo irreparable damage to his internal organs. According to reports coming from the World Health Organization (WHO), deficiencies of micronutrients are widespread in developing and underdeveloped areas where the facilities to perform advanced analyses are either not accessible or simply unavailable [1]. People living in such areas are often faced with difficulties such as poverty, absence of healthcare facilities, and a lack of understanding that result in delayed identification and treatment. Currently, the gold standard of analysis to identify vitamin deficiencies is laboratory analysis of the bloodstream. Though the results obtained are not inaccurate, the method of analyzing the bloodstream is both expensive and reliant on infrastructure and expertise. Therefore, people living a long way from such areas may be forced to either delay the test or simply not to undertake it at all. One of the first observations in nutrition medicine is that many vitamin deficiencies are clearly reflected on exterior parts of the body, including the skin, eyes, tongue, nails, lips, and hair. For example, vitamin A deficiency can lead to a condition causing eye dryness and night blindness; vitamin B-complex deficiencies can lead to a range of signs involving the tongue and brittle nails, while iron and zinc deficiencies tend to affect skin and hair [2].

These visible signs therefore provide an opportunity to study non-invasive diagnostic methods based on a visual examination rather than biochemical analysis. Recent accelerated advancements in the field of artificial intelligence (AI) led to the evolution of deep learning, an efficient paradigm for analyzing images, especially within the domain of medicine. Out of several deep learning models, Convolutional Neural Networks (CNN) were shown to possess a unique capability for deep learning tasks, wherein they could automatically learn discriminative features from raw images without even considering hand-crafted extraction of features

[3]. Traditional CNN architectures, including VGG models, were instrumental in paving the way for image classification tasks at a larger scale through the use of deep levels of visual features [4]. Later, more superior models like Residual Networks (ResNet) were developed with the help of skip connections to overcome the challenges of the vanishing gradient problem, allowing the network to learn features correctly even with deep levels [5, 6]. Recently, transfer learning has also hastened the implementation process for deep learning methods in healthcare applications. By using knowledge gained from large-scale datasets like ImageNet, models can then be fine-tuned for medical imaging applications when there is limited labeled data available, thereby speeding up training processes and improving generalization performance [7]. Notably, promising performance has also been realized through the use of deep learning models for various medical applications. For example, deep models have demonstrated dermatologist-level performance for classifying skin conditions [8], as well as radiologist-level performance for analyzing chest X-ray scans [9]. Various survey articles have also ascertained the general application of deep learning approaches for medical image classification, detection, or segmentation applications [10]. However, despite such advances, the task of identifying vitamin deficiencies given images is a difficult challenge. This is mainly because the signs of such deficiencies are often subtle, varied, and vary greatly from one person to another depending on factors like skin tone, lighting, age, and co-morbidities.

In addition, the features that are of primary interest to the identification of such deficiencies may be restricted to a limited part of the image, such as the conjunctiva of the eye or certain portions of the tongue. In such cases, standard CNN models may fail to highlight such important regions. To counter this shortcoming, attention-based localizing methods called Class Activation Mapping (CAM) [11], as well as Gradient-weighted Class Activation Mapping (Grad-CAM) [12], have been proposed to show areas of importance for making predictions on an image. In addition to improving classification accuracy, these approaches increase model explainability, which remains very important for developing trustworthy AI-supported healthcare systems [13]. Explainable Artificial Intelligence (XAI) enables model decision-making rationales to be interpreted by medical professionals or end-users, a requirement indispensable for responsible AI application in medicine [14]. However, another significant aspect of AI models designed for the medical field is computational complexity. In many practical applications of the healthcare industry, the computational requirement of the deploying device, such as a cell phone or a resource-limited embedded system, is a serious concern. However, some models like EfficientNet [15] and MobileNetV2 [16] have been designed to provide the best of both worlds: accuracy and efficiency. Moreover, recent studies on models such as Vision Transformer (ViT) have shown interesting results for learning global interactions of the image through the operation of self-attention, outperforming other models on certain medical imaging tasks [17]. Methods like model compression and knowledge distillation can be employed to reduce the burden of resource constraints on deploying AI models [18]. In light of the above findings, this study aims at the design of a vitamin deficiency diagnostic model based on the use of artificial intelligence through the use of deep learning approaches. The proposed approach uses images of visible human body parts in the diagnosis of vitamin deficiencies based on patterns created through the use of attention-based approaches in vitamin deficiency diagnosis. The proposed solution aims at the early awareness-creation process of vitamin deficiencies with respect to the primary function of its design, which is not intended for clinical diagnosis. In this regard, the proposed solution seeks to meet the goals of preventive health programs in underserved areas due to the unavailability of health diagnostic services.

II. Literature Survey

Initially, the area of medical image analysis focused on conventional image processing methods, as well as manually designed extraction techniques, that were sensitive to changes in image quality and lacking generalization power. However, the advent of deep learning impacted this area substantially. Deep residual learning has been proposed in [1], tackling the issue of gradient disappearance, which enabled the construction of extremely deep convolutional neural networks, bringing significant improvements to image classification tasks. Following work verified the efficiency of deep learning models for the task of challenging visual recognition, as reviewed in [2]. Currently, there is a lot of interest in deep learning applications in medical imaging, in which many diagnostic capabilities have been investigated. Esteva et al. [3] achieved dermatologist-level accuracy in skin cancer classification using deep neural networks, proving that even photographic images can contain clinically meaningful diagnostic information. Thus, developing the approach to the interpretability of CNN-driven medical diagnostics systems, Selvaraju et al. [4] presented Grad-CAM, focusing the visualization on those important image regions that influence model predictions. Later, Dimauro et al. [5] continued the development of image-based diagnosis with the detection of anemia from digital images of the palpebral conjunctiva. These examples further prove that at least in certain body areas, subtle visual appearances may signal health conditions. Rajpurkar et al. [6] also demonstrated the capabilities of deep learning methods to attain expert-level performance in chest radiograph diagnosis, again illustrating clinical applicability that now opens pathways for AI-based image analysis. Attention mechanism emerged as a crucial development in medical image processing. Xie et al. presented an extensive survey on attention mechanism-based methods that mentioned utilizing attention for improving model performance on critical areas. For efficiency with high precision, Tan & Le proposed EfficientNet models, and Ronneberger et al. suggested the use of U-Net in biomedical image segmentation that gained relevance in medical image processing tasks. Additionally, recent breakthroughs have also delved into modeling global context with attention mechanisms through transformer models. Dosovitskiy et al. [10] proposed the Vision Transformers model to address long-range dependencies by applying self-attention mechanisms effectively. In fact, the prior CNN models like VGG nets created by Simonyan and Zisserman [11] as well as the principles of deep learning described by Goodfellow et al. [12] served as a basis for developing such breakthroughs.

Also, applying transfer learning to model adaptation to specific medical domains with fewer labeled samples was emphasized by Pan and Yang [13]. Lightweight and optimized architectures are proposed to support the deployment in resource-constrained environments. Sandler et al. [14] introduced MobileNetV2 for efficient deep learning, and developments on residual learning by He et al. [15] furthermore strengthened the use of deep backbones when performing transfer learning. From a health application perspective, the explainability of AI has gained great importance; for instance, Samek et al. [16] highlighted the need for transparent and interpretable deep learning systems to support trust from clinicians. Apart from CNNs, region-based learning approaches such as R-CNN introduced by Girshick et al. [17] and class activation mapping techniques proposed by Zhou et al. [18] have improved localization and interpretability for visual recognition tasks. For the purpose of making it feasible to run on low-resource systems, distillation methods developed by Hinton et al. [19] to compress deep models without losing much accuracy have been considered.

In terms of public health, The World Health Organization emphasized that micronutrient deficiencies are still an ongoing concern in global healthcare, especially in developing areas, where analysis facilities are not readily available [20]. Current studies focusing on healthcare practices, in particular, reinforced that CNN models can efficiently contribute to making diagnoses of medical conditions based on visual biomarkers [21], while feature selection methods driven by optimization can ensure reliability in predictions [22], [23]. Deep learning with optimization methods was also tested in classifying neurologic diseases [24], analyzing residual images [25]. Although considerable advancements in AI-assisted medical diagnosis have been observed, existing studies only deal with a single medical issue or a body region. So far, very few studies have attempted the diagnosis of a vitamin deficiency in multi-regions based on visible symptoms. Therefore, the new proposed approach aims to combine ResNet152V2 with transfer learning and the attention module in order to examine the eyes, nails, skin, tongue, and hair to create a non-invasive vitamin deficiency diagnostic system.

III. METHODOLOGY

In this section, a brief overview of the approach and design of the proposed AI-assisted vitamin deficiency detection system is being provided. The approach is based on a non-invasive detection technique developed for identifying vitamin deficiencies based on observable symptoms present in external body parts of individuals, specifically in the eyes, skin, nails, tongue, and hair. The proposed detection model is developed utilizing deep learning methods, where transfer learning with ResNet152V2 is used for accurate feature detection, and the attention tool is also appended for improved detection of vitamin-deficiency-related subtle patterns in images. The process not only deals with model development but also covers each phase of data preparation, model architecture development, training, testing, and implementation for efficient detection of vitamin deficiencies.

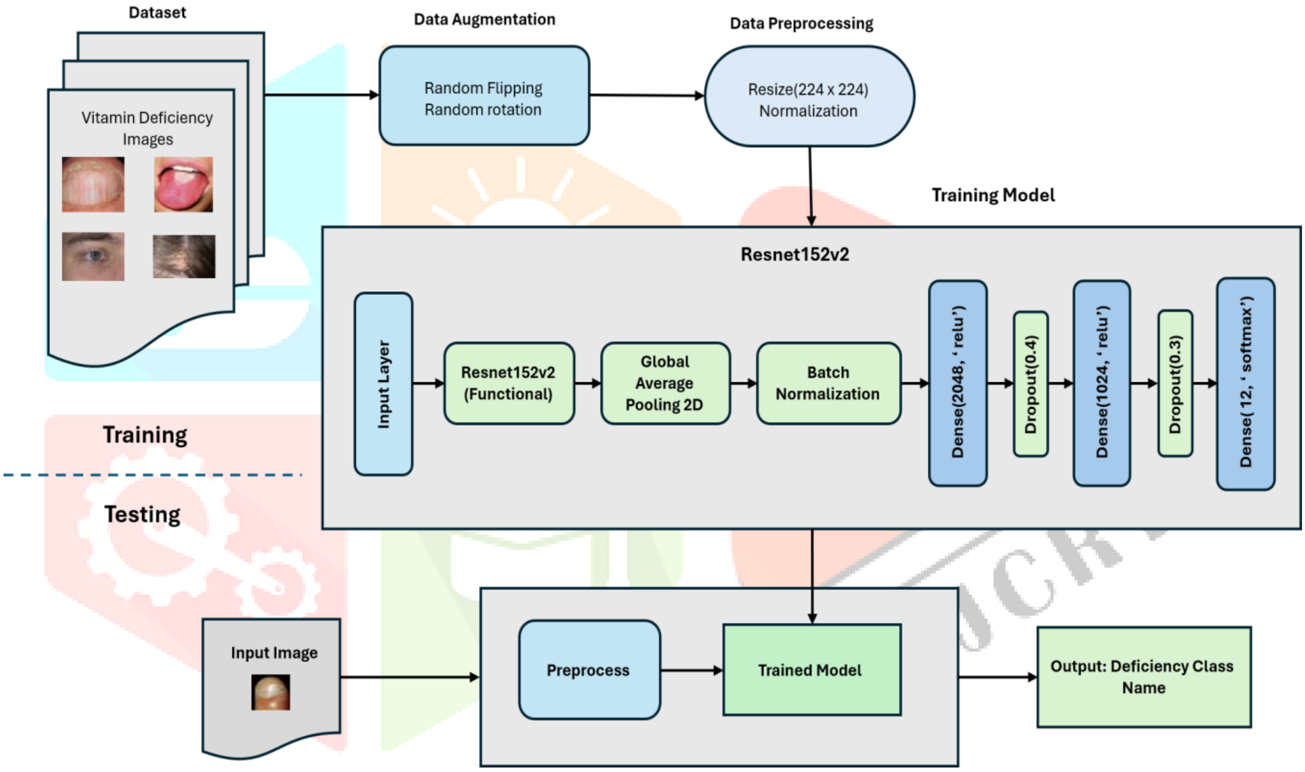

Fig 1: Overview of Resnet-based Framework for Vitamin Deficiency Detection

A. Dataset Preprocessing

The proposed method starts with the acquisition of images that demonstrate visible symptoms linked to deficiencies of vitamins and nutrients. The acquired images relate to body surfaces like the eyes, skin, nails, tongue, and hair because the visible symptoms like disc coloration of the tongue, brittle and pale nails, dry and red skin due to irritation, and hair changes can be easily noticed. The visible symptoms may include disc coloration of the tongue, brittle and pale nails, dry and red skin due to irritation, and hair changes. The images can be acquired from medical collections or uploaded by users on a website or a mobile phone.For overcoming some difficulties such as ill availability of data and unbalanced classes, in addition to variance in image orientation, data augmentation methods are employed in training. Some methods involve random flipping and rotating, where both methods are implemented artificially to increase variability among the dataset and also provide invariance in visual features to the system when training to correctly predict images taken from different angles in reality. As such, after augmentations, all images are resized to a fixed resolution of 224 × 224 pixels because it is the standard required input size for the ResNet152V2 architecture. Resizing also makes all samples uniform and saves much computational overhead. Besides, image normalization is done to scale the pixel intensity values in the same range. It improves numerical stability, speeds up convergence during training, and hence smooths the gradient updates in deep neural networks.

B. Feature Extraction Using ResNet152V2

The preprocessed images are then fed into the ResNet152V2 convolution neural network, which is considered the main feature extractor in the proposed system. ResNet152V2 is identified as one of the deep residual networks that comprises 152 layers and makes use of the idea of residuals to skip some layers and thus overcome the problem that leads to vanishing gradients during training. Transfer learning strategy is used in this work. Initially, the ResNet152V2 model is used, and its weights are pretrained on the ImageNet dataset. As such, the pretrained model utilizes knowledge from low- to mid-level vision learned from the ImageNet dataset, such as edges, gradients, textures, and shapes, to carry out the classification task. These representations are later fine-tuned to adapt to the deficiency representations in the medical images. Such a deep process of feature extraction helps the network learn high-level representations that capture subtle patterns associated with vitamin deficiencies, including color variations on the tongue, skin texture irregularities, structural changes in nails, and visual abnormalities in hair and eyes. The extracted feature maps give a compact yet informative representation of the input image for subsequent classification.

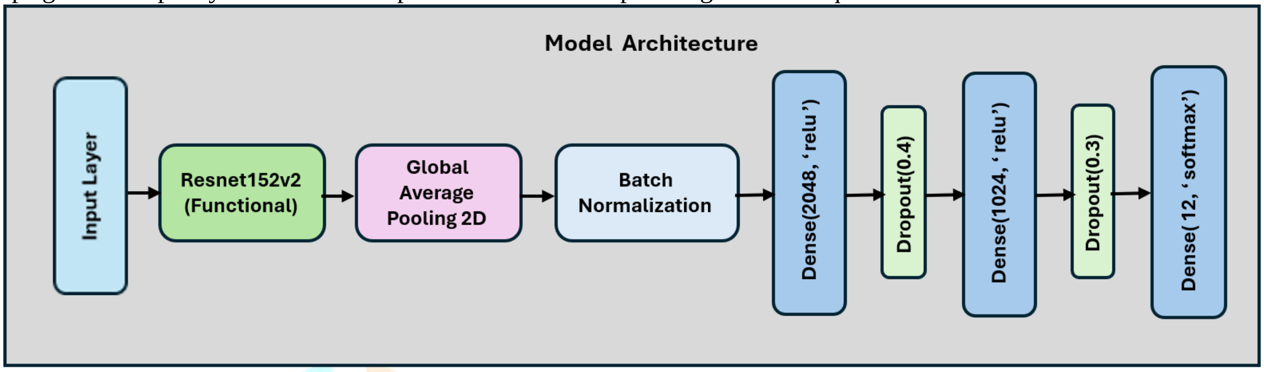

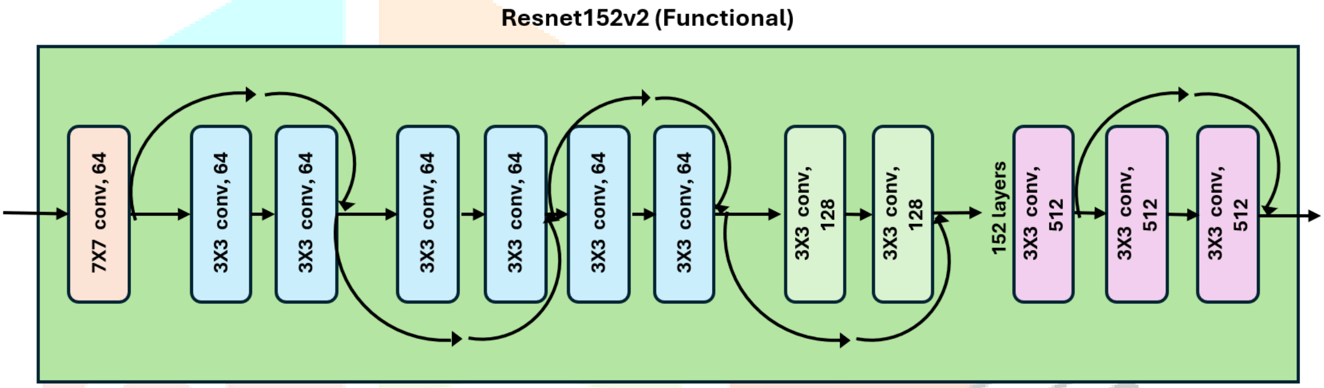

Fig.2: Proposed model architecture based on a ResNet152v2 backbone with global average pooling and fully connected classification layers

C. Classification and Model Training

After feature extraction, the feature maps undergo processing via a technique called Global Average Pooling (GAP). GAP allows for a reduction in spatial dimensions of feature maps while capturing key global features, thus reducing model parameters and overfitting associated with smaller medical image datasets when dealing with deep architectures. In addition to stabilizing the training process, a technique for faster convergence is applied, namely Batch Normalization. Batch normalization helps normalize internal activation values, reducing internal covariate shift so that a larger learning rate can be used. A custom classification head is then incorporated into this network architecture. This classification head contains several fully connected Dense layers that include ReLU activation functions to introduce nonlinearity into the feature spaces to allow for complex decision boundaries to be learned during training. Dropout layers are also included between each Dense layer in this custom classification head to randomly drop neural network connections during training to prevent features from co-adapting in an attempt to avoid overfitting in the classification head's performance for generalization. The final Dense layer contains a Softmax activation function to provide a probability score for each vitamin/micronutrient deficiency class. Training is conducted using the Adam optimizer, which immediately benefits from the adaptability that results from the use of learning rates, along with the benefit of momentum in the process of optimization, thereby achieving faster convergence. The categorical cross-entropy loss function is applied to calculate the difference that exists between the estimated probabilities and the actual class labels. Training is done in such a way that both the training metrics as well as the validation metrics are being watched. Furthermore, methods such as Early Stopping are used.

D. Prediction and Output Generation

In the inference step, the new image uploaded by the user is also subjected to the same pre-processing configurations that were done during training. The image is then passed through the trained ResNet152V2 model for classification. The features and classification process is done in one step. The output of this process is the predicted class of vitamin or micronutrient deficiency, as well as the confidence level of this prediction. By providing such a useful output result in a user-friendly format, it is quite easier to understand this process result without the need to understand technical details. The result can provide users with the benefit of having prior knowledge of any possible deficiency. The proposed system is intended to be a supportive, non-invasive screening tool for diagnosis, as opposed to a substitute for professional medical diagnosis. The proposed system will be highly automated, making it amenable for utilization in a resource-constrained environment where access to blood laboratory work is poor, thereby enabling early detection, preventive health, and improved health outcomes. This solution will greatly enable screening, prevention, early detection, and improved health via preventive health strategies based on AI-driven diagnostic technologies.

IV. EXPERIMENTAL SETUP

To assess the experimental validation of the proposed system meant for vitamin deficiency identification, a supervised deep learning model has been used. Here, the image dataset includes labeled pictures of different vitamin deficiencies like Vitamin A, Vitamin B complex, Vitamin C, Vitamin D, Vitamin E, Vitamin K, a class related to combined micronutrient deficiencies. These pictures are taken of the visible body parts like the eyes, skin, tongue, nails, lips, and hair. All the images are already preprocessed by resizing them into a fixed size of 224 x 224 pixels with normalized pixels. The data has been divided into training, validation, and testing for proper evaluation of efficiency. The ResNet152V2 model pretrained with the ImageNet dataset was used as the base model for transfer learning. The original classification layer of the model was modified and replaced with a set of fully connected layers based on the number of deficiency classes. The softmax function was applied for the final prediction. The model was trained with the Adam optimizer and the category cross-entropy loss function. The common evaluation metrics for the model were accuracy, precision, recall, and F1 score. The confusion matrix analysis for the model evaluation was conducted.

V. RESULTS AND DISCUSSIONS

A. Learning Behaviour and Convergence Analysis

The learning behavior of the ResNet152V2-based vitamin deficiency detection model was analyzed using training and validation loss and accuracy curves. The training loss decreased rapidly over epochs, while the validation loss closely followed this trend, indicating good generalization to unseen data. Minor variations in validation loss during later epochs are common in biomedical image analysis and were effectively controlled using dropout and early stopping. Both training and validation accuracy showed a steady increase and converged to high values, with a minimal gap between them, demonstrating that the model efficiently learns robust visual features for accurate vitamin deficiency detection.

The graph named “Training vs. Validation Accuracy” shows the accuracy level of the model for a number of epochs. At the beginning, the accuracy of both training and validation is low. This is because the model is still learning. During the training process, there is a rapid increase in accuracy. This indicates that the model has been able to train on the discriminative features associated with various vitamin deficiencies. The training accuracy keeps rising to a high value, while the validation accuracy tracks a similar pattern with slight oscillations. The small difference between the training accuracy and the validation accuracy during the training phase shows how well the model is doing with no serious overfitting.

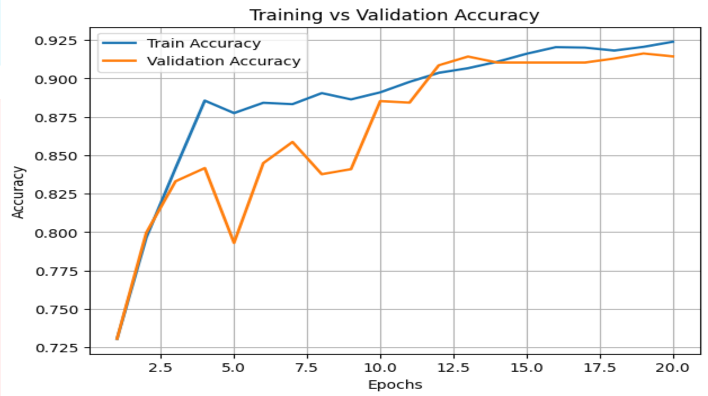

Fig 3: Training vs Validation Accuracy Graph

A minor drop in validation accuracy for the intermediate epochs can be explained by the difficulty of recognizing patterns in multi-class vitamin deficiencies and, at the same time, the presence of similar visual symptoms from different deficiency classes. Nevertheless, the validation accuracy recovers and stabilizes at a high level in later epochs, confirming the robustness of the trained model. The fact that both curves converge to higher accuracy values demonstrates the ResNet152V2 backbone has captured high-level visual representations, thus leading to efficient and reliable classification of vitamin and micronutrient deficiencies from visible symptoms.

"Training vs. Validation Loss" shows the change in the value of losses of the proposed ResNet152V2 model over the training epochs. When the model starts to train (in Epoch 1), it is expected that the value of the training loss will be high due to the random variables used to initialize model parameters. As the model continues its training, it shows a steep drop in the value of the training loss in the first few epochs, which reveals that it efficiently extracts relevant visual information with regards to the cases of vitamin deficiency. The validation loss is also steadily dropping, tracing the same pattern as the train loss. This is a clear indication that the network is not just memorizing the data in the train dataset but is actually learning a representation of the data that generalizes well on unseen data points in the validation dataset. This is because both curves are tracking very similarly.

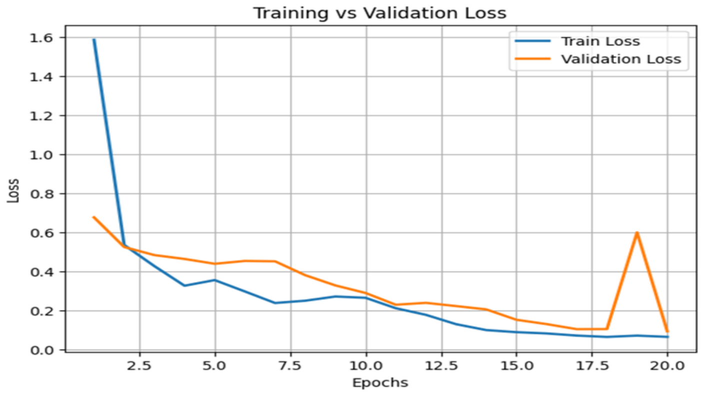

Fig 4: Training vs Validation Loss Graph

There is a slight variation in the validation loss in later epochs, which is typical in deep learning models when training images are from the actual medical setting, having various levels of quality, illumination, and presenting various symptoms. However, the graph validates that the validation loss converges to a low value, thereby authenticating the efficiency of the methods used for regularization, such as dropout, batch normalization, and early stopping. There is no issue in the convergence of the loss.

B. Quantitative Performance Evaluation and Metric Analysis

To assess the efficiency of the proposed ResNet152V2-based system for detection of Vitamin deficiency, the regular multi-classification evaluation criteria of Accuracy, Precision, Recall, and F1-score are used. The need for these criteria exists in the present project as the proposed system shall be utilized as a non-invasive tool for deficiency detection of Vitamins/Micronutrients based on the deficiency symptoms visible in the eye, skin, nails, tongue, and hair of the health concerns. In health care applications like the present project, the consequence of incorrect detection, especially the type of deficiency missed, gets affected by delayed treatments for better health outcomes.

The following terms will be defined as part of the vitamin deficiency classification task:

- TP (True Positives): Images that are labeled as particular vitamin/mineral deficiency when actually true.
- TN (True Negatives): Images that are correctly labeled as non-deficient or other classes of deficiency.
- FP (False Positives): These images are identified as a possible deficiency.
- FN (False Negatives): Images of true cases of deficiency incorrectly predicted to be non-deficient.

Accuracy is the level of correctness of the designed system in the classification process of vitamin deficiency images. If the accuracy is high, then it means that the system is able to properly classify the images into various categories of deficiency as well as normal images. However, in practical medical scanning systems, accuracy is an inadequate measure.

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

Precision is an assessment of the accuracy of deficiency predictions made by the system. Regarding vitamin deficiency, high precision ensures that if it predicts, say, Vitamin B2 or Vitamin B12 deficiency, it is most probable that it is correct. As it reduces false positives, there is less cause to fret and, hence, further tests might not be deemed necessary in healthy individuals.

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

Recall measures the capability of the system to identify actual deficiency cases correctly. Being a very important aspect of this project, if the vitamin deficiency is missed (i.e., false negative), it can lead to late treatment. A high recall value ensures the proposed model is properly identifying deficiency-related visual symptoms present at various body locations.

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

The F1-score gives a well-rounded assessment to the proposed system through the integration of the precision and recall scores. This metric works well in the current project since its implementation ensures good prediction performance as well as a strong ability to correctly identify deficiency cases. A higher score on the F1 measure shows good handling between avoiding false alerts and avoiding missing a diagnosis.

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

The proposed ResNet152V2-based vitamin deficiency detection system achieved an overall test accuracy of 91.4%, demonstrating strong generalization performance on unseen data. The model attained a precision of 91.0%, recall of 90.7%, and an F1-score of 90.8%, indicating a balanced trade-off between reliable deficiency prediction and effective identification of actual deficiency cases. These results confirm the suitability of the proposed system as a non-invasive screening and decision-support tool for vitamin deficiency detection.

Table 1.Overall Performance Evaluation

Metric	Accuracy	Precision	Recall	F1- Score
Value	91.4	91	90.7	90.8

C. System Deployment and Real-Time Prediction Analysis

This proposed system for the detection of vitamin deficiencies is implemented in the form of a user-friendly web application that allows the user to carry out the detection in a real-time, contactless manner. As part of the implementation, the developed ResNet152V2 deep neural network is coupled in the system's frontend. Here, the user is able to log into the system securely and upload pictures of the affected areas in the body, such as the skin, eyes, nails, tongue, or hair. These pictures will be automatically pre-processed in the same way as in the training process.The model, after preprocessing, extracts features and classes based on these features for predicting a possible vitamin/micronutrient deficiency for a corresponding image. The result is presented along with the image that has been uploaded, thereby making it easy for a user to understand a result. The model is also user-friendly, so it can be applied for massive screening, especially in a remote area where there is a lack of resources. The model, though not a diagnosis tool, is an assistance tool in health that emphasizes preventive health.

Fig 5: Home page

In this figure, it is possible to observe the landing page of the final system that was deployed to identify vitamin deficiency. The design of the landing page clearly illustrates the use of the system with technologies such as AI image identification, among others. The design is clear and easy to use. The UI focuses on early intervention and preventive health care, where it allows patients to interact with the system with little hassle. This allows it to be easily deployed in real-world settings, especially where mass screenings occur.

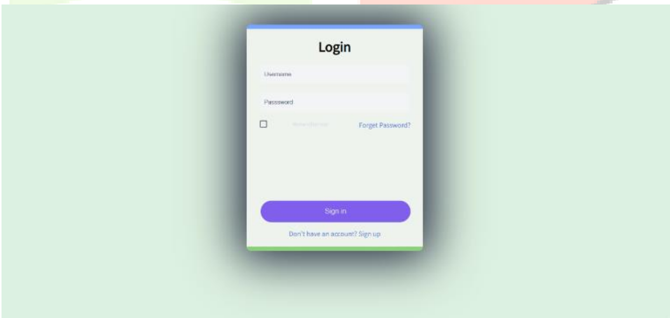

Fig 6: Login Page

It shows the login interface of the system, which allows authorized persons to securely log in. It also has an authentication system for controlling usage as well as protecting critical prediction information. Such a simple login interface is quite useful, especially when basic security is involved. The same interface also makes the system scalable by allowing its extension in the future to serve specific roles, whether health providers, administrative staff, or patients, hence improving the system's dependability.

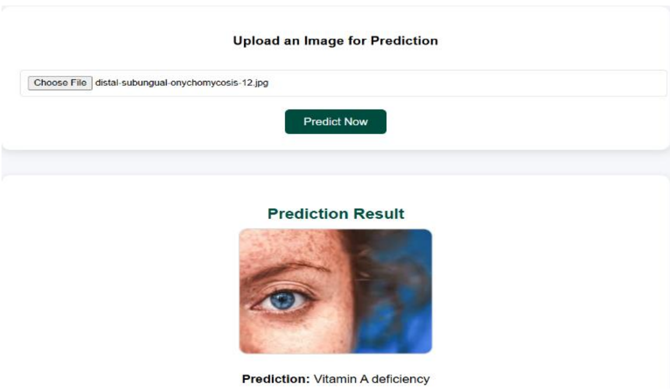

Fig 7:Prediction Result

In this figure, one can observe the prediction result produced using the trained ResNet152V2 model. The uploaded image is presented side by side with the result of thevitamin deficiency class prediction, indicating the class of vitamin deficiency, for instance, Vitamin A deficiency. The result displayed is an indication that the system is capable of doing predictive work in real time. The result interface also supports the fact that the system is a helpful tool for diagnosis and that it is capable of doing its work fast and without being invasive.

IV. CONCLUSION

The project introduced a framework for the non-invasive diagnosis of vitamin and nutrient deficiencies by observing visible signs on the external parts of the body, including the skin, eyes, tongue, hair, and nails. Using the ResNet V2-152 deep learning model and transfer learning, the project successfully trained on the distinctive characteristics in the visual signs to determine the different kinds of deficiencies. The combination of efficient preprocessing and regularizers helped in the successful training. The experimental outcome shows the excellent result of the proposed approach on the classification task, with a testing accuracy of 91.4 percent and very good precision, recall, and F1-measure. The near perfect similarity of the curves shows the absence of overfitting and supports the effectiveness of the approach towards generalizing well. The usage of the trained approach as a web application supports the practical viability of the approach, as predictions are possible with a very simple interface. In conclusion, the systems developed are useful support systems that can also be employed during the early stages of screening and awareness programs regarding vitamin deficiencies, especially in resource-poor areas that lack the express facility of laboratory testing. Although the aim is clearly not a substitute for diagnostic systems, the potential role of image analysis based on deep learning algorithms in preventive health practices has been highlighted.

REFERENCES

[1] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR), 2016, pp. 770–778.

[2] G. Litjens, T. Kooi, B. Ehteshami Bejnordi, et al., “A survey on deep learning in medical image analysis,” Medical Image Analysis, vol. 42, pp. 60–88, 2017.

[3] A. Esteva, B. Kuprel, R. A. Novoa, et al., “Dermatologist-level classification of skin cancer with deep neural networks,” Nature, vol. 542, pp. 115–118, 2017.

[4] R. R. Selvaraju, M. Cogswell, A. Das, et al., “Grad-CAM: Visual explanations from deep networks via gradient-based localization,” in Proc. IEEE/CVF International Conference on Computer Vision (ICCV), 2017, pp. 618–626.

[5] G. Dimauro, A. Guarini, D. Caivano, F. Girardi, and C. Pasquale, “Detecting clinical signs of anemia from digital images of the palpebral conjunctiva,” IEEE Access, vol. 7, pp. 113488–113498, 2019.

[6] P. Rajpurkar, J. Irvin, A. Bagul, et al., “Deep learning for chest radiograph diagnosis: A retrospective comparison of the CheXNeXt algorithm,” PLOS Medicine, vol. 15, no. 11, 2018.

[7] J. Xie, R. Liu, J. Luttrell, and C. Zhang, “Attention mechanisms in medical image analysis: A survey,” IEEE Access, vol. 9, pp. 7811–7830, 2021.

[8] M. Tan and Q. V. Le, “EfficientNet: Rethinking model scaling for convolutional neural networks,” in Proc. 36th Int. Conf. Machine Learning (ICML), 2019 (arXiv preprint).

[9] O. Ronneberger, P. Fischer, and T. Brox, “U-Net: Convolutional networks for biomedical image segmentation,” in Medical Image Computing and Computer-Assisted Intervention (MICCAI), 2015, pp. 234–241.

[10] A. Dosovitskiy, L. Beyer, A. Kolesnikov, et al., “An image is worth 16x16 words: Transformers for image recognition at scale” (Vision Transformer), 2020 (arXiv / ICLR).

[11] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” arXiv preprint (VGG nets).

[12] M. A. Goodfellow, Y. Bengio, and A. Courville, Deep Learning, MIT Press, 2016.

[13] S. Pan and Q. Yang, “A survey on transfer learning,” IEEE Transactions on Knowledge and Data Engineering, vol. 22, no. 10, pp. 1345–1359, 2010.

[14] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L. Chen, “MobileNetV2: Inverted residuals and linear bottlenecks,” in Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition (CVPR), 2018, pp. 4510–4520.

[15] K. He, X. Zhang, S. Ren, and J. Sun (ResNet followups) — (useful for transfer learning + backbone discussions). CVPR, 2016.

[16] S. Samek, M. T. Lu, and W. Samek, “Explainable AI: Interpreting, explaining and visualizing deep learning,” IEEE Signal Processing Magazine, 2019.

[17] R. Girshick, J. Donahue, T. Darrell, and J. Malik, “Rich feature hierarchies for accurate object detection and semantic segmentation (R-CNN),” in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2014.

[18] J. Zhou, O. Bau, S. Zoph, et al., “Learning deep features for discriminative localization” (CAM), in Proc. CVPR, 2016.

[19] G. Hinton, O. Vinyals, and J. Dean, “Distilling the knowledge in a neural network,” 2015 (arXiv), useful for model compression and edge deployment.

[20] World Health Organization, “Micronutrient deficiencies,” WHO Fact sheet, 2020.

[21] K. Guttikonda, Y. Ashvitha, V. S. R. Reddy, R. M. Krishna, and P. Sandeep, “Integrating convolutional neural networks and machine learning for accurate identification of autism spectrum disorder using facial biomarkers,” in Proc. IEEE Int. Conf. Emerging Systems and Intelligent Computing, pp. 1–6, 2024.

[22] K. Guttikonda, G. Ramachandran, and G. V. S. N. R. V. Prasad, “Autism spectrum disorder prediction using LASSO regularised bat search optimisation,” Int. J. Serv. Oper. Informatics, vol. 13, no. 1, pp. 1–20, 2024.

[23] K. Guttikonda G. Ramachandran, and G. V. S. N. R. V. Prasad, “Cuckoo search optimisation-based feature selection for predicting autism spectrum disorder using artificial immune algorithms,” J. Theor. Appl. Inf. Technol., vol. 103, no. 2, pp. 421–432, 2025.

[24] Apparna Allada, Rajaram Bhavani, Kavitha Chaduvula, Rajaram Priya,“Alzheimer's disease classification using competitive swarm multi-verse optimizer-based deep neuro-fuzzy network”, published in the journal of Concurrency and Computation: Practice and Experience, Vol.35, Issue 21, 25, e7696, September 2023.

[25] D.N.V.S.L.S. Indira, Babu Rao Markapudi, Kavitha Chaduvula, Rathna Jyothi Chaduvula, “Visual and buying sequence features-based product image recommendation using optimization based deep residual network”, published in the journal of Gene Expression Patterns, Volume 45, September 2022, PP: 119261.