

DEEP LEARNING-BASED EARLY DETECTION OF MALNUTRITION IN CHILDREN UNDER FIVE YEARS

DR. PRASANTHI LAKSHMI MALYALA¹, ALLAM AKSHITHA², SABBANI RISHITHA³, SWARAJ MAHENDRA LAKHE⁴, MUCHINTHALA KARTHIK⁵

ASSOCIATE PROFESSOR¹, UG SCHOLAR^{2,3,4&5}

DEPARTMENT OF AI&ML, NARSIMHA REDDY ENGINEERING COLLEGE (UGC-AUTONOMOUS) MAISAMMAGUDA (V), KOMPALLY, SECUNDERABAD, TELANGANA-500100

ABSTRACT

Malnutrition remains a critical public health concern, particularly among children, necessitating early and accurate diagnostic methods to support timely intervention. This project presents a deep learning-based system for the automatic detection of malnutrition in children using full-body pediatric images. The system classifies images into two categories—Malnutrition and Nutrition—leveraging Convolutional Neural Network architectures. The application is developed using Python for backend processing and Flask as the web framework, with a responsive interface built using HTML, CSS, and JavaScript. Two state-of-the-art deep learning models, ResNet152V2 and MobileNet, were trained and evaluated independently on a balanced dataset comprising 2,010 images (1,005 for each class). The image preprocessing pipeline includes rescaling, resizing, data augmentation, directory-based loading, batching, shuffling, and training-validation splitting, ensuring robust training and model generalization. The ResNet152V2 model achieved a training accuracy of 99% and test accuracy of 96%, while the MobileNet model recorded a training accuracy of 99% and test accuracy of 95%. For performance evaluation, standard classification metrics such as Precision, Recall, F1-score, and Confusion Matrix were utilized. Additionally, model accuracy and loss graphs were plotted to visualize training performance across epochs for both models. The results demonstrate that deep learning can effectively distinguish between malnourished and nourished children using visual cues alone, offering a scalable and non-invasive approach to nutritional screening in clinical and rural settings. This system has the potential to support healthcare professionals and aid in the early detection of malnutrition, thereby improving pediatric health outcomes.

INTRODUCTION

Malnutrition in children remains one of the most pressing global health challenges, particularly in low- and middle-income countries. It significantly impairs physical and cognitive development, weakens immunity, and increases the risk of morbidity and mortality among children under five. Traditional malnutrition assessment methods such as anthropometric

measurements (e.g., weight-for-height, mid-upper arm circumference) often require trained personnel and medical tools, which may not be readily available in remote or underserved regions. In this context, the integration of artificial intelligence (AI) and computer vision techniques offers a promising, scalable alternative for early detection and intervention. This project aims to develop an automated system for detecting malnutrition in children using deep learning and image classification techniques. Unlike prior research that focuses on facial image analysis, this system utilizes full-body pediatric images to classify nutritional status into two distinct categories: Malnutrition and Nutrition. The classification is performed using two powerful convolutional neural network (CNN) models—ResNet152V2 and MobileNet—trained independently to ensure optimal performance. The entire system is built using Python as the core programming language, with a Flask web framework powering the backend and a responsive user interface developed using HTML, CSS, and JavaScript. Image preprocessing includes rescaling, resizing, data augmentation, batching, and shuffling to enhance model robustness and generalization. Evaluation is conducted using well-established performance metrics including accuracy, precision, recall, F1-score, and confusion matrix, supported by graphical visualization of model loss and accuracy. Through this approach, the project demonstrates the potential of deep learning to facilitate rapid, non-invasive, and accessible malnutrition screening, with the goal of supporting public health systems and aiding early intervention in pediatric care.

LITERATURE SURVEY

1) Uncovering the complexity of childhood under-nutrition through strain-level analysis of the gut microbiome

AUTHORS: L. Wang et al

Undernutrition (UN) is a critical public health issue that threatens the lives of children under five in developing countries. While evidence indicates the crucial role of the gut microbiome (GM) in UN pathogenesis, the strain-level inspection and bacterial co-occurrence network investigation in the GM of UN children are lacking. This study examines the strain compositions of the GM in 61 undernutrition patients (UN group) and 36 healthy children (HC group) and explores the topological features of GM co-

occurrence networks using a complex network strategy. The strain-level annotation reveals that the differentially enriched species between the UN and HC groups are due to discriminated strain compositions. For example, *Prevotella copri* is mainly composed of *P. copri* ASM1680343v1 and *P. copri* ASM345920v1 in the HC group, but it is composed of *P. copri* ASM346549v1 and *P. copri* ASM347465v1 in the UN group. In addition, the UN-risk model constructed at the strain level demonstrates higher accuracy (AUC = 0.810) than that at the species level (AUC = 0.743). With complex network analysis, we further discovered that the UN group had a more complex GM co-occurrence network, with more hub bacteria and a higher clustering coefficient but lower information transfer efficiencies. Moreover, the results at the strain level suggested the inaccurate and even false conclusions obtained from species level analysis. Overall, this study highlights the importance of examining the GM at the strain level and investigating bacterial co-occurrence networks to advance our knowledge of UN pathogenesis.

2) Identification of malnutrition and prediction of BMI from facial images using real-time image processing and machine learning

AUTHORS: C. Dhanamjayulu et al

Human faces contain useful information that can be used in the identification of age, gender, weight etc. Among these biometrics, body mass index (BMI) and body weight are good indicators of a healthy person. Motivated by the recent health science studies, this work investigates ways to identify malnutrition affected people and obese people by analyzing body weight and BMI from facial images by proposing a regression method based on the 50-layers Residual network architecture. For face detection, Multi-task Cascaded Convolutional Neural Networks have been employed. A system is created to evaluate BMI along with age and gender from human facial real-time images. Malnutrition and obesity are commonly determined with the help of BMI. In the previous works, height, weight, and BMI estimation through automatic means have predominantly focused on full-body images and videos of humans. The usage of facial images for estimating such traits have been given less importance. In order to facilitate the analysis, the dataset is cleaned along with metadata containing information about the persons height, weight, age, and gender. Gender-based analysis is performed for the prediction of BMI. Finally, an email containing the persons picture along with their details is sent to the concerned health officer.

3) Establishing a machine learning model for predicting nutritional risk through facial feature recognition

AUTHORS: J. Wang, C. He, and Z. Long

Malnutrition affects many worldwide, necessitating accurate and timely nutritional risk assessment. This study aims to develop and validate a machine learning model using facial feature recognition for predicting nutritional risk. This innovative approach seeks to offer a non-invasive, efficient method for early identification and intervention, ultimately improving health outcomes. We gathered medical examination data and facial images from 949 patients across multiple hospitals to predict nutritional status. In this multicenter investigation, facial images underwent preprocessing via face alignment and cropping. Orbital fat pads were isolated using the U-net model, with the histogram of oriented gradient (HOG) method employed for feature extraction. Standardized HOG features were subjected to principal component analysis (PCA) for dimensionality reduction. A support vector machine (SVM) classification model was utilized for NRS-2002 detection. Our approach established a non-linear mapping between facial features and NRS-2002 nutritional risk scores, providing an innovative method for evaluating patient nutritional status. In context of orbital fat pad area segmentation with U-net model, the averaged dice coefficient is 88.3%. Our experimental results show that the proposed method to predict NRS-2002 scores achieves an accuracy of 73.1%. We also grouped the samples by gender, age, and the location of the hospital where the data were collected to evaluate the classification accuracy in different subsets. The classification accuracy rate for the elderly group was 85%, while the non-elderly group exhibited a classification accuracy rate of 71.1%; Furthermore, the classification accuracy rate for males and females were 69.2 and 78.6%, respectively. Hospitals located in remote areas, such as Tibet and Yunnan, yielded a classification accuracy rate of 76.5% for collected patient samples, whereas hospitals in non-remote areas achieved a classification accuracy rate of 71.1%.

4) A Deep Learning Approach for Malnutrition Detection

AUTHORS: S. Ankalaki, V. G. Biradar, P. Kishor Kumar Naik, and G. S. Hukkeri

The timely detection of malnutrition in children is of paramount importance, as it allows for early intervention and treatment. This proactive approach not only prevents further health deterioration but also fosters proper growth, minimizing the long-term consequences of malnutrition, such as stunted growth, impaired cognitive development, and increased vulnerability to diseases. Our work encompasses the creation of a new dataset comprising images of children in Healthy, Undernourished, Stunting, and Wasting categories. The core objective is to assess the deep learning model performance in classifying these children images. The experimentation is carried out by varying epochs, batch size, optimizers AdamW, Adamax, and RMSprop; and different values of the learning rate 0.1, 0.01, 0.001, and 0.0001 during model

training. The model is trained on image dataset constructed by cleaning images generated by the stable diffusion model. The model is tested on randomly selected child images from websites. The model successfully classified two classes with 95% accuracy, 97.6% F1 score, precision 97.6%, and 97.6% recall with Adam optimizers, 0.0001 learning rate, and Batch size 4. Additionally, for the four-class categorization scenario, the study broadens the classification. The model achieved 88.87% accuracy, 90.3% recall, 90.2% precision, and an F1 score of 90% for four-class categorization with AdamW optimization, 0.0001 learning rate, and batch size 6. These results are satisfactory for prediction of malnutrition category in children.

5) Artificial intelligence for the practical assessment of nutritional status in emergencies

AUTHORS: B. Watkins, L. Odallo, and J. Yu

This paper describes a novel method for detecting child malnutrition based on artificial intelligence and facial photography. Estimates of severe and moderate acute malnutrition in children are critical for rapid emergency responses. However, the two traditional measurement methods, mid-upper arm circumference (MUAC) and weight-for-height (WFH), are impractical in conflict and catastrophic disaster situations. They require well-trained enumerators, cumbersome equipment, and close supervision. The Method for Extremely Rapid Observation of Nutritional Status (MERON) addresses the problem, using simple facial photographs. Facial features are extracted to predict Body Mass Index (BMI) in adults and Weight for Height Z Score (WFHZ) in children under five. MERON correctly predicts adult BMI classification with 78% accuracy. A variant of the model, trained on a sample of 3167 children in Kenya, successfully classified 60% of cases. On most measures, MERON was easier and more culturally acceptable to use than the traditional measurement methods. If MERON were to be trained and validated on a larger sample, with more extreme cases, it would provide a practical solution to a recurrent humanitarian problem.

SYSTEM ANALYSIS

EXISTING SYSTEM:

- ❖ The existing system titled "Detection of Malnutrition in Children Under 5 Years of Old Using Deep Learning" presents a deep learning-based approach for classifying the nutritional status of children using image data. The system primarily focuses on classifying children into three categories: Normal, At Risk of Malnutrition, and Severely Malnourished, based on visual features extracted from images.

- ❖ The existing system methodology follows the CRISP-DM (Cross-Industry Standard Process for Data Mining) framework, which involves structured stages such as business understanding, data preparation, modeling, and evaluation.
- ❖ The dataset used in the existing system consists of over 1,000 images sourced from the Cuna Más National Program, the Open Images Dataset V7, and other open sources. Preprocessing is conducted using MTCNN (Multi-task Cascaded Convolutional Networks) for face detection, followed by image resizing, horizontal flipping, zooming, and face extraction to enhance model robustness.
- ❖ The existing system evaluates four Convolutional Neural Network (CNN) models: ResNet50, EfficientNet-B4, VGG16, and AlexNet. These models are trained separately, and their performances are assessed using metrics such as accuracy, precision, recall, and F1-score. Among these, ResNet50 achieved the highest accuracy of 92%, proving most effective in identifying malnutrition patterns from images.
- ❖ Although the existing system focused on model performance evaluation, it also proposes the future integration of the best-performing model into a mobile application for real-world deployment. The application architecture includes modules for nutritional analysis, user data management, educational content, and personalized recommendations, with a frontend built using Flutter and a backend hosted on AWS.
- ❖ This system serves as a comprehensive research effort in applying AI and deep learning to address pediatric malnutrition, combining CNN-based classification with practical considerations for field deployment.

DISADVANTAGES OF EXISTING SYSTEM:

- ❖ While the existing system demonstrates the potential of deep learning in detecting malnutrition among children under five, it presents several limitations that impact its practical implementation and scalability. One of the primary disadvantages is the reliance on facial images for classification, which may not always be readily available or ethically feasible to collect, especially in under-resourced or remote areas. Furthermore, facial features alone may not sufficiently capture the physical symptoms of malnutrition that

appear in other parts of the body, potentially affecting diagnostic accuracy.

- ❖ The system classifies images into three categories—Normal, At Risk, and Severely Malnourished—which may introduce subjectivity in borderline cases. In real-world clinical settings, such nuanced distinctions can be challenging to validate without anthropometric data, and may lead to misclassification, particularly when image quality varies due to environmental conditions such as lighting, background clutter, or child posture.
- ❖ Another drawback is the limited dataset size (just over 1,000 images), which may not adequately represent the diversity of malnutrition cases across different regions, ethnicities, and environmental backgrounds. This constraint can hinder the model's generalization ability when exposed to new, unseen data from varied populations.
- ❖ In addition, while the system identifies ResNet50 as the most effective model with 92% accuracy, it does not address model deployment optimization, such as reducing computational load or ensuring compatibility with low-power mobile devices. The absence of real-world deployment and testing also limits the validation of the model's robustness in uncontrolled environments.
- ❖ Lastly, though a mobile application is proposed, it remains conceptual within the study. The system lacks a fully integrated, end-to-end solution that combines the trained model with user-facing features such as real-time image capture, offline functionality, and clinical feedback mechanisms—essential components for widespread adoption in rural healthcare settings.

PROPOSED SYSTEM:

- ❖ The proposed system introduces a deep learning-based image classification framework designed to detect malnutrition in children using full-body pediatric images. The system focuses on classification, categorizing input images into two classes: Malnutrition and Nutrition. Unlike facial image-based systems, this approach utilizes full-body images to capture more comprehensive visual indicators of nutritional status.
- ❖ The core functionality is implemented using the Python programming language, with a Flask web framework to manage server-side operations and HTML, CSS, and JavaScript for the front-end

interface. This setup enables users to upload images through a web-based application for instant classification.

- ❖ Two deep convolutional neural network (CNN) models—ResNet152V2 and MobileNet—are trained independently on a dataset consisting of 2,010 full-body images, evenly distributed between the two classes (1,005 images each for Malnutrition and Nutrition). The image preprocessing pipeline includes steps such as image rescaling, resizing, data augmentation, batching, shuffling, and training-validation splitting to ensure model consistency and robustness.
- ❖ During training, both models are evaluated using standard performance metrics including Precision, Recall, F1-score, and Confusion Matrix to assess classification effectiveness. Additionally, Model Accuracy and Loss graphs are generated and visualized to monitor the training progress and model behavior over time.
- ❖ The system is deployed through a web interface that accepts image inputs and returns classification results in real-time. Each model's predictions are handled independently, allowing comparative analysis and modular integration for future enhancements.

ADVANTAGES OF PROPOSED SYSTEM:

- ❖ The proposed system offers a number of significant advantages over traditional and existing methods for malnutrition detection. By utilizing full-body images instead of relying solely on facial images or anthropometric measurements, the system captures a broader range of visual indicators associated with nutritional health, leading to more accurate classification.
- ❖ The use of advanced deep learning models such as ResNet152V2 and MobileNet enables high-performance image classification, with both models achieving high training and testing accuracy. ResNet152V2, with its deeper architecture, effectively extracts complex visual features, while MobileNet ensures efficient performance with lightweight deployment potential. This dual-model setup allows flexibility based on application constraints like computational resources.
- ❖ The system's preprocessing pipeline—which includes image augmentation, resizing, rescaling, and shuffling—enhances generalization and reduces

overfitting, leading to reliable model behavior across diverse input samples. Additionally, the balanced dataset used for training ensures fair learning across both malnutrition and nutrition classes, minimizing bias.

- ❖ From a usability perspective, the integration of the system into a Flask-based web application with a user-friendly front-end allows healthcare workers, researchers, or field staff to upload images and receive instant classification results. This contributes to real-time, remote, and non-invasive diagnosis, which is especially valuable in rural or resource-limited settings where clinical equipment or personnel may not be readily available.
- ❖ Furthermore, performance analysis using metrics such as Precision, Recall, F1-score, and Confusion Matrix, along with training loss and accuracy visualizations, provides transparency and interpretability of model performance. This supports ongoing model evaluation and potential tuning for improved reliability.
- ❖ In summary, the proposed system combines technical robustness with practical deployment capability, offering a scalable and accessible solution for early detection and monitoring of malnutrition in children.


Dataset Source	Cuna Más, Open Images V7, Online Resources (1,000+ images)	Custom dataset with 2,010 full-body images (1,005 malnutrition + 1,005 nutrition)
Deployment	Mobile application proposed (conceptual only)	Fully developed web-based system using Flask, HTML, CSS, JavaScript
Evaluation Metrics	Accuracy, Precision, Recall, F1-score	Accuracy, Precision, Recall, F1-score, Confusion Matrix
Performance Visualization	Confusion matrix only	Confusion Matrix, Model Loss Graph, Model Accuracy Graph
User Interaction	Not implemented	Real-time image upload and prediction via web interface
Framework Used	No specific implementation framework mentioned	Flask Web Framework with Python Backend

IMPLEMENTATION

SYSTEM ARCHITECTURE:

Comparison Between Existing System and Proposed System

Aspect	Existing System	Proposed System
System Title	Detection of Malnutrition in Children Under 5 Years of Old Using Deep Learning	Detection of Malnutrition in Children Through Visual Diagnosis of Pediatric Images Using Deep Learning
Image Type Used	Facial Images	Full-Body Images
Number of Classes	Three Classes: Normal, At Risk of Malnutrition, Severely Malnourished	Two Classes: Malnutrition and Nutrition
CNN Models Used	ResNet50, EfficientNet-B4, AlexNet, VGG16	ResNet152V2, MobileNet
Best Model Accuracy	ResNet50 – 92%	ResNet152V2 – 96% Test Accuracy MobileNet – 95% Test Accuracy
Preprocessing Techniques	Face detection (MTCNN), resizing, flipping, zooming	Image rescaling, resizing, data augmentation, batching, shuffling, train/validation split



MODULES:

Data Collection:

- ❖ In the first module of the Malnutrition detection using deep learning, we make the data collection process.

This is the first real step towards the real development of a deep learning model, collecting data. This is a critical step that will cascade in how good the model will be, the more and better data that we get; the better our model will perform.

- ❖ There are several techniques to collect the data, like web scraping, manual interventions. The dataset is located in the model folder. The dataset is referred from the popular dataset repository called kaggle. The following is the link of the dataset:

- ❖ Kaggle Dataset Link:

<https://www.kaggle.com/datasets/jayaprakashpondy/children-nutrition>

Dataset:

- ❖ In this module, Set up two main directories: training, and testing. Training, and testing directories, create subdirectories is 2 class label.
- ❖ Place the images into their respective class directories based on their labels. Each image should be placed in the directory corresponding to its class label and by organizing the dataset in this structure, it becomes straightforward to load the data into your deep learning framework for training and testing purposes and Total dataset size is 2010.

Data Preparation:

- ❖ During the data preparation stage, it is crucial to preprocess the data to ensure it is suitable for training. This involves tasks such as resizing images to a standard size, normalizing pixel values, and encoding labels if necessary. To achieve this, the ImageDataGenerator from Keras can be utilized. For instance, to resize images to a standard size of 224x224 pixels, the Target_size parameter can be set to (img_height, img_width) = (224, 224).
- ❖ Additionally, pixel values can be normalized by setting the rescale parameter to 1./255, which scales the pixel values to be between 0 and 1. Furthermore, data augmentation techniques such as random shear and zoom can be applied to enhance the training data. By leveraging these techniques, the data can be effectively preprocessed to improve the performance of the deep learning model.

Feature Extraction:

- ❖ For models like ResNet152V2 and MobileNet, which come pre-trained with feature extraction layers,

explicit feature extraction may not always be necessary. By setting trainable = False, we freeze these pre-trained layers, allowing them to retain their learned representations while preventing further updates during training.

- ❖ In the context of ResNet152V2 and MobileNet, setting trainable = False ensures that the weights of the feature extraction layers remain fixed during training. This approach is commonly used in transfer learning scenarios, where the pre-trained model is fine-tuned on a new dataset for a specific task, such as image classification or object detection.
- ❖ By adopting this strategy, we strike a balance between leveraging powerful pre-trained representations and adapting the model to our specific dataset, ultimately improving both training efficiency and model performance.

Splitting the dataset:

- ❖ Divide your dataset into training and validation to evaluate your model's performance. Typically, you might use an 80-20 split, but this can vary based on your dataset size and specific requirements.

Model Selection:

- ❖ The training module is responsible for training the deep learning models using the preprocessed data. It implements two popular architectures: ResNet152V2 and MobileNet.

ResNet152V2:

- ❖ ResNet152V2 is structured with a series of residual blocks, each containing multiple convolutional layers. The architecture starts with a single convolutional layer followed by a batch normalization layer and a ReLU activation. This is followed by a series of residual blocks, grouped into stages with increasing numbers of filters. The network concludes with a global average pooling layer and a fully connected layer that outputs the final classification.

MobileNet:

- ❖ MobileNet models are composed of multiple layers of depthwise separable convolutions, structured in a modular fashion. The architecture typically begins with a standard convolution layer followed by a series of depthwise separable convolutions. Each layer is

followed by batch normalization and ReLU activation. This sequence of operations efficiently extracts features while maintaining a low computational footprint

Training the Model:

- ❖ To train the models, the training module first loads the pre-trained weights for ResNet152V2 and MobileNet from the ImageNet dataset. It then adds a global average pooling layer and a fully connected layer with the appropriate number of classes for the specific task. The base layers of the pre-trained models are frozen by setting their trainable flag to False, allowing only the added layers to be trained.
- ❖ The training process involves optimizing the model parameters using a suitable optimization algorithm, such as Adam or SGD, and a loss function appropriate for the task (e.g., categorical cross-entropy for classification). The training data is fed to the model in batches, and the gradients are computed and used to update the model weights. The training process continues for a specified number of epochs or until a certain performance metric is achieved on a validation set.
- ❖ After training, the module saves the trained models for future use in the prediction and evaluation modules. The saved models can be loaded and used for inference on new data or fine-tuned on additional datasets if needed.

Analyze and Prediction:

- ❖ Once training is complete, analyze the training process (e.g., loss curves) and make predictions on your validation set to assess the model's performance.

Accuracy on test set:

- ❖ Once the model is trained, it needs to be evaluated for its performance. This module involves splitting the dataset into training and testing subsets and assessing the model's accuracy, precision, recall, and F1-score.
- ❖ The ResNet152V2 architecture achieves a Training accuracy of 99.00% and validation accuracy of 95.00%. The MobileNet architecture attains a Training accuracy of 99.00% and validation accuracy of 95.00%

Saving the Trained Model:

- ❖ Once you're confident enough to take your trained and tested model into the production-ready environment, the first step is to save it into an .h5 or .pkl file using a library like pickle.
- ❖ Make sure you have pickle installed in your environment.
- ❖ Next, let's import the module and dump the model into .pkl file.

Prediction Module:

- ❖ Develop a prediction module to make predictions using the trained ResNet152V2 and MobileNet model. This module should take input images, preprocess them as necessary, and output predictions

Model Evaluation Module

- ❖ This module evaluates the performance of the trained models using the testing dataset. It calculates accuracy metrics and other performance indicators to assess model effectiveness.
- ❖ Evaluate model accuracy, precision, recall, and F1-score.
- ❖ Generate confusion matrices for both models.
- ❖ Compare the performance of the ResNet152V2 and MobileNet.
- ❖ Accuracy, precision, recall, and F1-score are used to evaluate model performance.

CONCLUSION

The project successfully demonstrates the feasibility and effectiveness of employing deep learning techniques for the automatic detection of malnutrition in children using full-body images. By leveraging advanced convolutional neural network architectures, namely ResNet152V2 and MobileNet, the system achieved high classification accuracy, with test scores reaching 96% and 95%, respectively. The comprehensive preprocessing pipeline—including image rescaling, resizing, augmentation, and structured data partitioning—contributed to model stability and enhanced predictive performance.

The integration of the trained models into a Flask-based web application allows real-time classification of pediatric images into Malnutrition and Nutrition categories, offering an accessible and user-friendly diagnostic tool. Evaluation metrics such as precision, recall, F1-score, and confusion matrices further validated the reliability and robustness of the system.

Overall, this project highlights the potential of deep learning-driven visual diagnostic systems to support early identification of malnutrition in children, providing a scalable, non-invasive, and efficient approach that can complement traditional assessment methods in clinical and community health settings.

FUTURE WORK:

While the current system demonstrates promising results in detecting malnutrition using full-body pediatric images, there are several directions in which this work can be extended. One key area is the expansion of the dataset to include a more diverse set of images across different ethnicities, age groups, lighting conditions, and geographic backgrounds to improve the model's generalization and robustness in real-world scenarios.

Incorporating additional classes such as "At Risk of Malnutrition" or "Severely Malnourished" can provide a more granular understanding of nutritional health, enabling better clinical decision-making. Furthermore, integrating anthropometric data and metadata such as age, weight, or region, alongside image data, can enrich the feature space and potentially increase diagnostic accuracy through multimodal learning.

Another potential enhancement is the development of a mobile application version of the system to facilitate easy and offline usage in remote or low-resource areas. Model optimization techniques such as quantization and pruning can be employed to ensure that deep learning models run efficiently on low-power devices.

Finally, implementing explainable AI (XAI) techniques would help interpret the model's decisions, thereby increasing trust and transparency among healthcare providers and end-users. Conducting clinical trials or field validation studies could also help assess the system's impact in practical healthcare settings, paving the way for its adoption in public health initiatives.

REFERENCES

- [1] L. Wang et al., "Uncovering the complexity of childhood undernutrition through strain-level analysis of the gut microbiome," *BMC Microbiol*, vol. 24, no. 1, Dec. 2024, doi: 10.1186/s12866-024-03211-w.
- [2] "World Health Organization." Accessed: Oct. 21, 2024. [Online]. Available: https://www.who.int/es/health-topics/childgrowth#tab=tab_1
- [3] "Instituto Nacional de Estadística e Informática." Accessed: Oct. 21, 2024. [Online]. Available:

<https://m.inei.gob.pe/prensa/noticias/el-431-de-la-poblacion-de-6-a-35-meses-de-edad-sufrio-de-anemia-en-el-ano-2023-15077/>

- [4] V. Malathy, N. Maiti, N. Kumar, D. Lavanya, S. Aswath, and S. B. Banu, "Deep Learning -Enhanced Image Segmentation for Medical Diagnostics," in 2024 International Conference on Advances in Computing, Communication and Applied Informatics (ACCAI), 2024, pp. 1–6. doi: 10.1109/ACCAI61061.2024.10602242.
- [5] C. Dhanamjayulu et al., "Identification of malnutrition and prediction of BMI from facial images using real-time image processing and machine learning," *IET Image Process*, vol. 16, no. 3, pp. 647–658, 2022, doi: 10.1049/ipr2.12222.
- [6] J. Wang, C. He, and Z. Long, "Establishing a machine learning model for predicting nutritional risk through facial feature recognition," *Front Nutr*, vol. 10, 2023, doi: 10.3389/fnut.2023.1219193.
- [7] S. Ankalaki, V. G. Biradar, P. Kishor Kumar Naik, and G. S. Hukkeri, "A Deep Learning Approach for Malnutrition Detection," *International journal of online and biomedical engineering*, vol. 20, no. 6, pp. 116–138, 2024, doi: 10.3991/ijoe.v20i06.46919.
- [8] B. Watkins, L. Odallo, and J. Yu, "Artificial intelligence for the practical assessment of nutritional status in emergencies," *Expert Syst*, 2024, doi: 10.1111/EXSY.13550.
- [9] A. Goel and S. Neduncheliyan, "Evaluating the Efficiency of Deep Learning Models for Personalized Healthcare," in 2023 IEEE Engineering Informatics, 2023, pp. 1–6. doi: 10.1109/IEEECONF58110.2023.10520611.
- [10] V. G. Biradar, H. H. P. K. Pareek, S. T. K. Aishwarya, and K. Aneeth, "Performance Analysis of Deep Learning Models for Detection of Malnutrition in Children," in 2024 International Conference on Data Science and Network Security (ICDSNS), 2024, pp. 1–8. doi: 10.1109/ICDSNS62112.2024.10691135.
- [11] C. Schröer, F. Kruse, and J. M. Gómez, "A Systematic Literature Review on Applying CRISP-DM Process Model," *Procedia Comput Sci*, vol. 181, pp. 526–534, Jan. 2021, doi: 10.1016/J.PROCS.2021.01.199.
- [12] "Programa Nacional Cuna Más." Accessed: Oct. 21, 2024. [Online]. Available:

<https://www.gob.pe/institucion/cunamas/noticias?filter%5Bterm%5D=&filter%5Btype%5D=&sheet=14>

[13] "Open Images Dataset V7." Accessed: Oct. 21, 2024. [Online]. Available: <https://storage.googleapis.com/openimages/web/visualizer/index.html?type=detection&set=train&c=%2Fm%2F01bl7v>

[14] J. Xiang and G. Zhu, "Joint face detection and facial expression recognition with MTCNN," in Proceedings - 2017 4th

International Conference on Information Science and Control Engineering, ICISCE 2017, Institute of Electrical and Electronics Engineers Inc., Nov. 2017, pp. 424–427. doi: 10.1109/ICISCE.2017.95.

[15] D. R. E. Rajappan, S. Sanith, and S. Subbulakshmi, "Malnutrition Detection using Deep Learning Models," in 2023 4th IEEE Global Conference for Advancement in Technology (GCAT), 2023, pp. 1–8. doi: 10.1109/GCAT59970.2023.10353327.