NutriAI: AI-Powered Child Malnutrition Assessment in Low-Resource Environments

Misaal Khan[1,2], **Shivang Agarwal**[1], **Mayank Vatsa**[1], **Richa Singh**[1] and **Kuldeep Singh**[2]
[1]Indian Institute of Technology Jodhpur, India
[2]All India Institute of Medical Science Jodhpur, India
{khan.9, shivangagarwal, mvatsa, richa}@iitj.ac.in, kulpra@gmail.com

Abstract

Malnutrition among infants and young children is a pervasive public health concern, particularly in developing countries where resources are limited. Millions of children globally suffer from malnourishment and its complications. Despite the best efforts of governments and organizations, malnourishment persists and remains a leading cause of morbidity and mortality among children under five. Physical measurements, such as weight, height, middle-upper-arm-circumference (muac), and head circumference are commonly used to assess the nutritional status of children. However, this approach can be resource-intensive and challenging to carry out on a large scale. In this research, we are developing NutriAI, a low-cost solution that leverages small sample size classification approach to detect malnutrition by analyzing 2D images of the subjects in multiple poses. The proposed solution will not only reduce the workload of health workers but also provide a more efficient means of monitoring the nutritional status of children. On the dataset prepared as part of this research, the baseline results highlight that the modern deep learning approaches can facilitate malnutrition detection via anthropometric indicators in the presence of diversity with respect to age, gender, physical characteristics, and accessories including clothing.

1 Problem Statement

Malnutrition affects millions of people globally and continues to pose a significant risk, despite advancements in medical science. It is responsible for nearly 45% of deaths among children under the age of five[1]. Malnutrition can manifest in various ways, including under-nutrition (wasting, stunting, underweight), insufficient vitamins or minerals, obesity, and resulting in diet-related noncommunicable diseases. Figure 1 shows sample images of healthy and malnourished children. In 2020, it was estimated that 149 million children under the age of five were stunted, 45 million were wasted, and 38.9 million were overweight or obese[2]. As evident from Figure 2, malnourishment is more prevalent in low and middle-income countries, indicating the alarming need to develop novel interventions. The global burden of malnutrition has severe and lasting effects on individuals, families, communities, and countries, with significant impacts on developmental, economic, social, and medical levels. Early nutrition is vital in providing the best start and long-term benefits.

Figure 1: Sample images from the internet curated malnourished images dataset, MalDB [Sethi *et al.*, 2020].

The most commonly used malnutrition screening tools are the WHO[3] and CDC's growth charts [Kuczmarski, 2002], which use weight for age (WFH) and height for age (HFA) indices to monitor and identify malnutrition in clinical and community settings by comparing children's growth and nutritional status. Additionally, clinical assessments may include examining for edema, hair thinning, a swollen abdomen, dry skin, and mineral/vitamin deficiencies [Ogden *et*

[1]https://www.who.int/news-room/fact-sheets/detail/malnutrition

[2]https://www.who.int/health-topics/malnutrition

[3]https://www.who.int/tools/child-growth-standards/standards

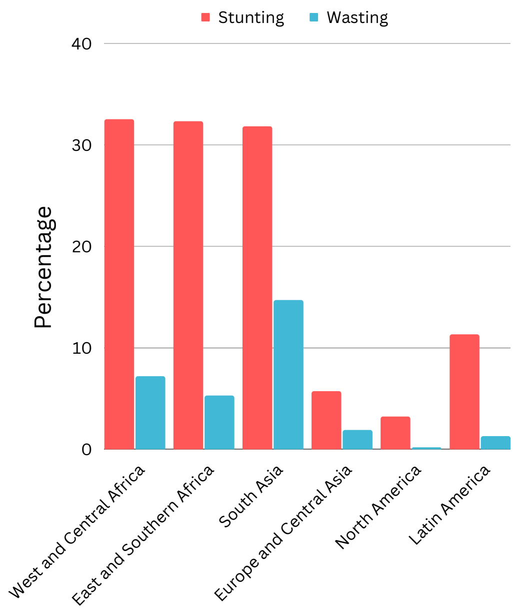

Figure 2: Percentage of children affected by stunting and wasting across various geographical locations. Source: UNICEF, WHO, World Bank Joint Child Malnutrition Estimates and country dataset, 2021 Edition.

al., 2002]. Currently, most of this process is manually executed and presents challenges with respect to measurement inaccuracy, variability, cultural factors, and sometimes inaccuracies of the measuring devices as well. Furthermore, these tools may be limited by using outdated growth benchmarks, which may not reflect modern population increases [Heemann *et al.*, 2021]. Specifically, the conventional methods have the following issues:

1. Subjectivity: The measurements taken by different healthcare workers may vary, leading to inconsistencies in the diagnosis of malnourishment.
2. Time-consuming: Measuring multiple body parts for each child can be time-consuming and may require multiple visits, which can be a burden for families.
3. Cost: Traditional methods may require specialized equipment and trained personnel, which can be expensive and not readily available in all settings.
4. Limited data: Traditional methods only provide limited data on body size and shape, which may not capture the full extent of malnourishment or other health conditions.

Computer vision and AI approaches can help overcome some of these limitations by providing a non-invasive, automated, and objective method to assess malnutrition [C *et al.*, 2022]. This paper focuses on developing NutriAI, AI-powered child malnutrition assessment system. The proposed system provides visible health assessment decision by analyzing images of a child's face and/or body and extract relevant features such as facial landmarks, body shape, and proportions. These models can increase measurement accuracy to customize and provide context-specific interventions.

2 Target SDGs and the Societal Impact

Young, undernourished children suffer from weakened immune systems, making them more susceptible to infections, mental development delays, long-term cognitive and psychological impairments, and stunted holistic development [Shrimpton *et al.*, 2001; Grantham-McGregor, 1995; Black *et al.*, 2008]. There is a malnutrition cycle in which dietary requirements are not reached in pregnant women in populations suffering from chronic malnutrition. As a result, babies delivered by these mothers have low birth weight, cannot achieve their full growth potential, and may be stunted, and vulnerable to diseases, sickness, and mortality at an early age. The cycle is exacerbated as low birth weight women evolve into undernourished children and adults, making them more prone to bear children with low birth weight [Dukhi, 2020].

NutriAI has the potential to improve health outcomes at both national and global levels by enabling detection of malnutrition in children at early stages, which can lead to appropriate treatment and nutritional support, reducing child morbidity and mortality rates in the country, and contributing to the achievement of Sustainable Development Goal - Newborn and child mortality (SDG 3.2). Moreover, the detection of malnutrition at early stages can prevent severe cases and associated health complications, thereby reducing healthcare costs, which is particularly relevant in low-resource settings. NutriAI can also provide valuable insights into the prevalence and distribution of malnutrition, enabling policymakers and healthcare providers to make data-driven decisions regarding national nutrition interventions and programs, and also contributing to the achievement of SDG 2.2 - Malnutrition.

On a global level, NutriAI can help improve equity in healthcare services, ensuring that all children have access to appropriate nutrition interventions, regardless of their location or socio-economic status. While SDG 2.2 and SDG 3.2 are the primary Sustainable Development Goals addressed by NutriAI, it can also contribute to achieving other related goals, such as improving access to safe, nutritious, and sufficient food (SDG 2.1) (by focusing nutritional interventions at the targeted population), reducing the number of deaths and illnesses caused by non-communicable diseases (SDG 3.4), and reducing inequalities by improving access to healthcare and reducing the burden of malnutrition, which disproportionately affects disadvantaged and marginalized communities (SDG 10.1).

3 Goals

Recent technological advancements, particularly the widespread use of smartphones equipped with high-quality cameras, have opened up new opportunities for nutritional screening [Khan *et al.*, 2022]. The aim is to develop an efficient system for nutritional status prediction that can automate the traditional anthropometric measurement based prediction methods at low costs while investigating the effectiveness of 2-D facial and full body images for prediction. This dual approach shall provide a more comprehensive and standardized screening tool for children's nutritional status. The primary hypothesis of this project involves several key points. Firstly, anthropometric body measurements are important for nutritional status screening. Secondly, the manual process of acquiring anthropometric measurements is time-consuming, labor-intensive, and prone to errors; therefore, the possibility of using facial and full body images must be explored.

The anticipated outcomes of this research include the exploration of the applicability of 2D images obtained from phone cameras in health screening without the use of additional sensors or training. The development of a continuous growth monitoring system based on the suggested model, and digitization of health parameters and records for better data management are also expected. Additionally, the target is to develop a more sustainable solution to reduce child mortality by assisting doctors in tracking milestones according to WHO standards of growth metrics. Finally, the ultimate outcome is developing a generalizable framework for addressing small sample-size problem statements in child healthcare using computer vision.

3.1 Technical Contributions

We propose a dual-branch framework for malnourishment screening that focuses on facial and full-body images. The first branch of the framework explores the feasibility of deep learning architectures for prediction from 2-D images. The second branch consists of a novel domain-information guided segmentation and regression for anthropometric measurement prediction. In the technical domain, this project explores novel few-shot [Wang *et al.*, 2020] learning approach with limited sample sizes. These proposed system incorporates domain information from medical experts, which improves model performance and provides explainability for the model's working, a common limitation in conventional black box models. The expected framework is generalizable to other small sample sizes and unexplored problem statements in child healthcare.

To address the scarcity of publicly available image datasets for malnourishment and nutritional status detection, one of the contributions of this project is to curate a sufficiently large dataset involving subjects of different ethnicity and their images in multiple poses. So far, we have collected $3,090$ images from the Internet and $2,729$ images pertaining to 334 subjects along with their anthropometric ground truths collected physically from government hospitals and schools. In addition, with the motivation of using transfer learning, we have collected a dataset of 80 adult subjects involving 479 images. The dataset will be made available for research purposes under a license agreement.

4 Proposed Framework

The NutriAI system proposed in this research aims to overcome the challenges imposed by traditional approaches and offer a scalable, trustworthy, and user-friendly framework to empower health workers in low-resource settings, reduce dependency on trained medical practitioners, and lower the cost of national health missions. In this section, we discuss the data collection protocol, baseline results of SOTA models, and the proposed framework.

4.1 Data Collection and Preparation

We utilized a smartphone camera, measuring tape, weighing machine, tripod, and laptop to collect data. Subjects provided full-body and facial images, captured under well-lit conditions, and wearing minimal or fitted clothing. Images were taken from frontal, lateral, and posterior angles, within a 3-5ft range. Additionally, facial images were obtained via selfies, maintaining a frontal face angle.

Anthropometric measurements recorded included height, middle-upper-arm circumference (MUAC), head circumference, and weight. Weight was recorded to the nearest gram, and all other measurements were in centimeters. We used a GUI application to record and consolidate these measurements and images into a dataset for model training and validation. Figure 4 represents a flowchart of the data collection protocol.

For accurate classification, ground truth labeling was conducted using World Health Organization (WHO) and Centers for Disease Control and Prevention (CDC) growth standards. This allowed us to categorize subjects into 'healthy' or 'undernourished' based on their z-scores. Clinical collaborators confirmed these labels [WHO, 1995].

Data collection is an ongoing effort; however, three datasets have been acquired so far. The first is an internet-curated set, with annotations by professional doctors. The second is from adult students at IIT Jodhpur, intended for transfer learning. Lastly, a clinical dataset was obtained from AIIMS Jodhpur and a nearby government school. Figure 5 shows some sample images from the proposed datasets. Details of the datasets are given below:

1. *MalDB [Sethi* et al.*, 2020]:* Comprising 3,090 images sourced from the internet, this dataset is annotated by professional doctors. The images, captured in uncontrolled settings, are equipped with binary labels indicating 'healthy' or 'not healthy'. Each category includes 1,545 samples. The skin color in the dataset is distributed as follows: 1,864 dark brown, 784 brown, 270 light brown, and 172 white skin images.
2. *IITJ (Adult):* This dataset, consisting of 479 images from 80 adult students at IITJ, was collected to enhance transfer learning for children. It comprises six different poses: Frontal, Back, Lateral left, Lateral right, Hands wide, and Facial selfies. Alongside the images, it contains the following anthropometric measurements: Height, Weight, MUAC, HC, Waist Circumference, Age, and Gender, as well as binary labels formulated based on the Body Mass Index (BMI).
3. *AIIMSJ and Govt. School (Child health dataset):* This dataset contains 2,729 images from 334 subjects, primarily children from clinical and community health settings. The dataset includes six distinct poses and similar anthropometric measurements as the IITJ dataset. Binary labels are provided based on the standards set by the WHO and the CDC.

4.2 Model Pipeline

The proposed system for classifying malnourishment is a dual-branch framework that combines conventional methods

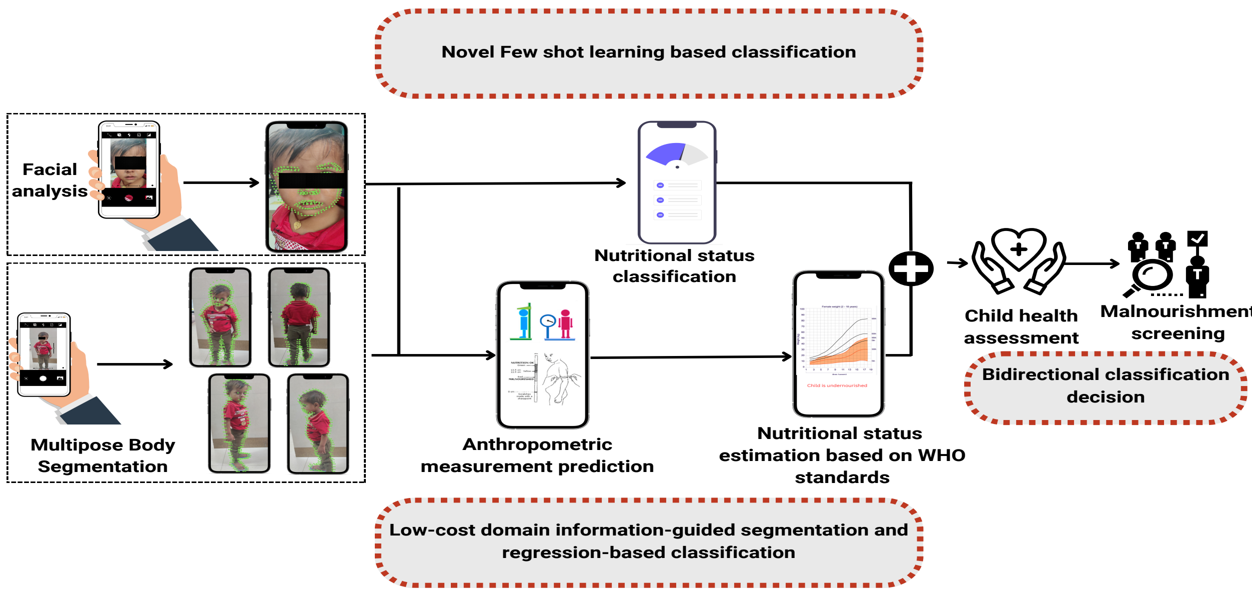

Figure 3: Multitask project pipeline of NutriAI: The proposed framework works on two parallel approaches: (1) the facial and full body images captured during data collection are passed to perform binary classification using few shot learning, and (2) the multi-pose body images are used to predict anthropometric measurements such as height, middle-upper-arm circumference (MUAC), head circumference, and waist circumference. These measurements are then compared with the growth standards developed by the World Health Organization (WHO) and Centers for Disease Control and Prevention (CDC) to predict the nutritional status of the child.

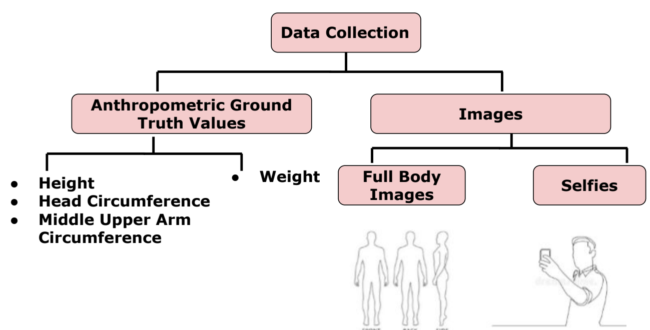

Figure 4: Data collection protocol

with AI-based pattern recognition. The two parallel branches work together to determine whether an individual is malnourished. In the first branch, we utilize few-shot learning classifiers to identify new facial and full-body features from images and predict malnourishment. The second branch uses novel domain information-guided segmentation and regression to classify healthy and not-healthy individuals based on accurate anthropometric measurements derived from facial and multi-pose full body images. This branch automates and improves conventional methods for greater efficiency and accuracy. By combining conventional and AI-based approaches, we offer a more dependable and generalizable screening system that can be trusted to provide accurate results. The multi-task network approach allows for a comprehensive evaluation of an individual's nutritional status.

Using two parallel pipelines, this multitask [Kendall *et al.*, 2018] approach in the NutriAI system allows for bidirectional validation of results. The system leverages the power of computer vision to provide more reliable and accurate screening decisions while also weighing in the contemporary approach of malnourishment screening, which allows for a more trustworthy model. Figure 3 is a pictorial representation of the proposed approach.

4.3 Baseline Results

The experimental protocol was designed to establish a baseline for binary classification accuracy using popular deep learning architectures and three proposed datasets: MalDB, IITJ (Adult), and Child Health Dataset. This baseline evaluation aimed to assess the performance of different models and provide a benchmark for future experiments involving more complex models. As discussed in the data collection and preparation section, the binary labels for the MalDB dataset were obtained through annotation by expert doctors, who determined whether each sample was classified as healthy or not healthy. For the IITJ dataset, the binary labels were obtained using BMI-based health prediction. The Full body image and bmi dataset of celebrities were also used for transfer learning to compensate small sample size of IITJ dataset. Lastly, the binary labels for the Child Health Dataset were derived through z-score evaluation from anthropometric mea-

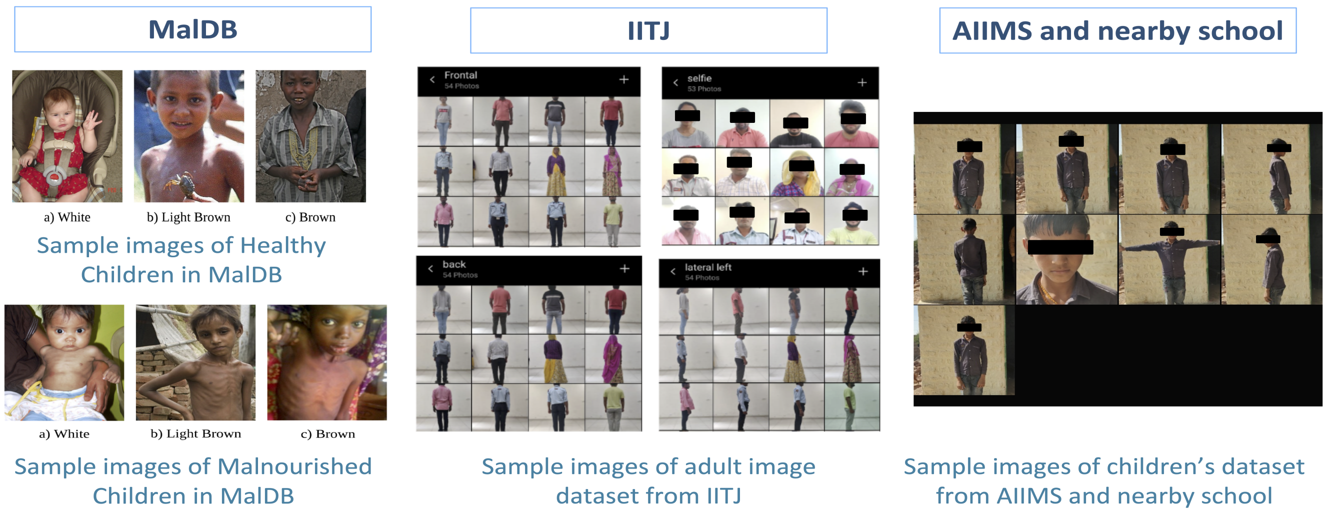

Figure 5: Some sample images from the curated datasets.

surements of the subjects, following the norms established by the World Health Organization (WHO) and the Centers for Disease Control and Prevention (CDC).The z-score represents the number of standard deviations that a particular measurement falls above or below the average value for a reference population. In the context of the Child Health Dataset, the anthropometric measurements of each subject were used to calculate z-scores for different attributes such as height-for-age, weight-for-age, and BMI-for-age. Based on these z-scores, binary labels were assigned to indicate the health status of the subjects. It is crucial to emphasize that the final classification agrees with the prognosis provided by expert doctors in the field.

We evaluated several widely recognized deep learning architectures, namely ResNet18, ResNet50, ResNet101 [He *et al.*, 2016], VGG16, VGG19 [Simonyan and Zisserman, 2015], and DenseNet [Huang *et al.*, 2017]. These architectures have demonstrated strong performance in various computer vision tasks.

To improve the performance of the models, we employed finetuning techniques by introducing three additional convolutional layers, each followed by ReLU activation, batch normalization, and dropout. The hyperparameter choices were made based on prior empirical evidence and served as the foundation for training the model on the given dataset with no. of training epochs - 10, batch size-32 and learning rate - 0.001 for Adam optimiser. To ensure fair evaluation, we employed a train-test-validation split of 60-20-20. The dataset was also augmented by applying transformations that do not alter the shape of individuals in the image to obtain the final accuracies. The experimental results presented in Table 1 are based on the dataset collected thus far, and it should be noted that this dataset will continue to be updated and expanded. The continuous updates and additions to the dataset will contribute to a more comprehensive evaluation of model performance.

It is worth noting that the ResNet models consistently outperformed the other architectures on all three datasets. ResNet achieved the highest accuracy on the MalDB and Child Health Dataset. However, VGG exhibited superior performance on the IITJ (Adult) dataset. Our implementation has been open-sourced for further research and is accessible online.[4]

Dataset	**ResNet**			**VGG**		**DenseNet**
	18	**50**	**101**	**16**	**19**	
MalDB	70.65	71.34	71.65	69.40	71.03	70.97
IITJ (Adult)	63.10	63.03	62.30	64.20	64.70	63.19
Child Health Dataset	63.20	67.34	63.66	68.5	67.78	65.50

Table 1: Baseline results on various datasets. The results demonstrates the model's performance on the task of binary classification, i.e., healthy and unhealthy.

4.4 Evaluation Criteria

To evaluate the performance of the NutriAI system, we compare its screening results with the ground truth labels established using the WHO and CDC growth standards and expert advise from doctors. Using standardized growth standards allows for comparing results across different populations and time periods and enhances the reliability and reproducibility of results. By combining the objective z-score evaluation with the doctors' professional assessment, the dataset's ground truth health labels reflect a comprehensive understanding of the subjects' health conditions. It also ensures that the model's performance is based on well-established and widely used guidelines for assessing nutritional status, and comparisons can be made with other studies using similar standards. Although we have demonstrated only accuracy in the baseline results, sensitivity, specificity, accuracy, F1 and

[4]Accessed on May 22, 2023. Available at: https://github.com/misaalkhan/ijcai_baselines

AUC-ROC scores are the desired evaluation metrics for the final NutriAI system for predicting healthy and undernourished children and comparing them with those obtained from the traditional screening methods.

5 Collaboration Opportunities and a Possible Roadmap

This proposal aims to reduce child mortality via AI techniques and a multidisciplinary approach, collaborating with child health experts from the All India Institute of Medical Science, Jodhpur (AIIMSJ), and AI professionals from the Indian Institute of Technology Jodhpur (IITJ). Additional input from nutrition, pediatrics, epidemiology, and public health scholars is welcome.

The strategy involves applying domain knowledge to deep learning's black-box methods, creating efficient tools for small sample-size issues. This approach will also enable continuous monitoring of child mortality rates and allow for policy recommendations on regional and international scales. We are working with domestic and international NGOs and Aasha workers (health and paramedical staff in India) for model testing and feedback, including an NGO in the aspirational district of Sirohi. This partnership facilitates data collection for model training and validation and offers insights into deploying the model in real-world settings.

The project comprises three tasks: (1) conceptualizing a framework for factors influencing child mortality rates; (2) developing the NutriAI model; and (3) deploying in public health care settings with NGO and ASHA worker collaboration. AI team members will develop NutriAI, while child health experts will focus on the conceptual framework and policy recommendations. Some of the outputs that the project will result in are as follows.

1. A white paper outlining child mortality and malnutrition in low-income countries, discussing NutriAI as a potential solution for targeted interventions, and addressing considerations for data collection, technical infrastructure, and stakeholder engagement.
2. Technical advancements in the field of AI: The algorithms developed as part of this project will serve as baseline to address similar problems not only in healthcare but in other similar domains as well. The datasets prepared as part of this project will be released to the research community to further the research on this important topic. We aim to publish the project outcomes in IJCAI or similar venue in 2024 or 2025.
3. The proposed tool NutriAI, will be deployed in primary healthcare settings in collaboration with ground workers. We are already working with ASHA workers and NGOs in two states of India, Punjab, and Rajasthan. By leveraging the existing infrastructure, we plan to connect with local administration for the deployment of NutriAI through networks established by NGOs and working closely with healthcare stakeholders to embed it within routine healthcare practices.

Expected Timeline: The project consists of several key stages spanning a two-year timeline. Its objective is to develop a computer vision-based model for predicting the nutritional status of children through facial and full-body images while also estimating anthropometric measurements and z-scores.

The initial phase of the project deals with data collection and analysis. A diverse dataset of children's images in multiple poses and their corresponding anthropometric data will be collected and preprocessed. Following the data analysis and image preprocessing, a segmentation algorithm will be developed to extract relevant features for further analysis. The subsequent phase involves model development and deployment. A computer vision-based model will be developed using the extracted features to predict anthropometric measurements and z-scores. A pilot deployment will also be carried out with the help of collaborating NGOs to simulate the model's performance in practical settings.

Continuous data augmentation will be implemented to improve the model's performance, contributing to the generalisation across various populations. Once the model development and data augmentation are complete, a final evaluation will be conducted. The project team will carefully document the methodology and code used throughout the various stages to ensure documentation and reproducibility. We plan to adhere to this timeline as closely as possible while allowing for flexibility and adjustments based on unforeseen circumstances or opportunities that may arise during the project.

6 Responsible Decision-Making in Malnourishment: Managing Risks and Challenges

Malnutrition is a complex issue with multiple underlying causes, which may not be fully accounted for by the model. Data scarcity in underprivileged regions may skew the model's accuracy. Furthermore, the model needs to account for variability in images such as skin color, gender, and age, as well as factors like lighting, pose, and background, which significantly affect performance.

Interpreting images to assess a child's nutritional status is a complex task requiring the model to understand the relationship between a child's appearance and their health. Addressing these challenges necessitates a large, diverse, high-quality image dataset, consistent monitoring, and model updates. Collaboration with healthcare professionals can enhance data collection and classification approaches. The model will address fairness concerning subjects' age, gender, and skin color, and safeguard data privacy and security.

Machine learning applications in diagnosing malnutrition raise ethical concerns. Accuracy is paramount to avoid misdiagnosis and potential harm. Issues of stigmatization and discrimination arise. Additionally, the lack of human interaction and empathy in the healthcare process when using a machine-learning model for diagnosis can potentially reduce the quality of care for children. Addressing these concerns calls for stakeholder engagement, including medical professionals, data scientists, and ethicists. Regulatory guidelines need to be established to safeguard personal information and curb potential harm. Routine performance checks help ensure accuracy and detect unintended consequences.

Finally, we plan to make the proposed methodology and code publicly available to promote reproducibility and facilitate the transferability of the approach to other contexts. The proposed solution is designed to be modular, allowing other researchers to customize and adapt the model to their specific needs and data. We also plan to conduct rigorous testing and evaluate the proposed solution to ensure its reproducibility and generalizability. We will provide detailed documentation of the experiments, including the datasets used, the training and validation procedures, and the evaluation metrics.

Ethical Justification

The proposed research project is grounded in the fundamental principles of medical ethics, as set out by the American Medical Association (AMA) [Riddick, 2003], which emphasize the importance of confidentiality, honesty, patient autonomy, human dignity, patient access to medical care, and responsibility to both the community and the patient. In addition, the study adheres to the National Ethical Guidelines for Biomedical and Health established by the Indian Council of Medical Research (ICMR) [Mathur *et al.*, 2018] and the Helsinki Declaration, modified in 2000 [Lie *et al.*, 2004] on health research considerations such as informed consent, privacy and confidentiality, vulnerable populations, independent ethical review, data sharing etc. To ensure compliance with ethical standards, the study requires all volunteers or their guardians, in the case of minors, to provide informed consent before participation. Participants are informed of their right to withdraw from the study at any time without providing a reason. The data collected ensures to not cause any immediate or long-term harm to the children. The anthropometric measurements are taken using a weighing machine and measuring tapes, while the full body and facial images are stored in a dataset that adheres to all privacy and security requirements.

Moreover, the study complies with the guidelines outlined by the World Health Organization (WHO) on the Ethics and Governance of Artificial Intelligence for Health [WHO, 2021] and the Royal Society's report on Machine Learning in Healthcare [Blacklaws, 2018], which advocates for the ethical use of AI in healthcare, emphasizing transparency, accountability, and the protection of privacy and confidentiality. We have obtained approval for the project from both AIIMS Jodhpur and IIT Jodhpur's ethics committees based on the same ethical justifications and project protocol. These approvals allow us to collect data in accordance with ethical guidelines and ensure that the study is conducted in an ethical and responsible manner.

Acknowledgments

Khan is partly supported through PMRF Fellowship and Vatsa is partially supported through SwarnaJayanti Fellowship by the Government of India. The authors also acknowledge Ashish Sethi for help in MalDB sample collection and doctors for discussion and for annotating some samples.

A Curricula Vitae

Misaal Khan is a Dual degree (Master-PhD) student in the joint program of Medical Technologies at All India Institute of Medical Sciences (AIIMS), Jodhpur, and Indian Institute of Technology (IIT), Jodhpur. An active part of the the Image Analysis and Biometrics Lab, IIT Jodhpur, her research interests lie in the fields of computer vision-aided medical diagnosis, few-shot learning, and explainable AI. Misaal has also received the prestigious Prime Minister's Research Fellowship. Before joining her PhD, she worked as a research assistant at the Indian Institute of Science (IISc), Bangalore where she worked on point-of-care disease diagnostics.

Shivang Agarwal is a Postdoctoral Fellow at the Image Analysis and Biometrics Lab, IIT Jodhpur. He received his PhD in Computer Science and Engineering from IIT (BHU), Varanasi, India. Prior to this, he obtained his Masters in computer science and engineering from NIT Rourkela. His research interests are face recognition, surveillance biometrics, and computer vision. He has published several papers in reputed journals and serves as a reviewer in top journals and conferences such as Pattern Recognition, Pattern Recognition Letters and NeurIPS.

Mayank Vatsa is a Professor at the Department of Computer Science and Engineering at the Indian Institute of Technology (IIT) Jodhpur, India. He serves as Dean of Research and Development and Project Director of the Technology and Innovation Hub (TIH) at IIT Jodhpur, India. He is a Fellow of IEEE and IAPR, and the recipient of the prestigious SwarnaJayanti Fellowship. He was a member of the Indian Biometrics Standards Committee for e-gov applications. He also served as a member of the UIDAI's Biometrics Standard Subcommittee member (face, fingerprint, and iris standards). He also served as the Vice President (Publications) for IEEE Biometrics Council, where he led the efforts to start IEEE Transactions on Biometrics, Behavior and Identity Science. His areas of interest are computer vision, machine learning, biometrics and Trusted AI.

Richa Singh is a Professor at the Department of CSE, IIT Jodhpur. She received Ph.D. degree in computer science from the West Virginia University, USA. She is a Fellow of IEEE and IAPR, and has received several awards including the Kusum and Mohandas Pai Faculty Research Fellowship at the IIIT-Delhi, the FAST Award by the Department of Science and Technology, India, and several best paper and best poster awards in international conferences. She has also held leadership positions in professional organizations, such as Program Co-Chair of CVPR2022, General Chair of FG2021, Vice President - Publications of the IEEE Biometrics Council, and an Associate Editor-in-Chief of Pattern Recognition.

Kuldeep Singh is currently serving as a Professor and Head of Pediatrics, as well as the Dean of Academics at All India Institute of Medical Science (AIIMS), Jodhpur. His expertise in child health, medical genetics, medical education, and public health is evident through his contributions to various projects such as the establishment of the National Inherited Administration Kendra (NIDAN) at AIIMS Jodhpur, GenomeIndia, and OneHealth Consortium. He also oversees the India-Sweden Innovation Center at AIIMS Jodhpur.

References

[Black *et al.*, 2008] Robert E Black, Lindsay H Allen, Zulfiqar A Bhutta, Laura E Caulfield, Mercedes De Onis, Majid Ezzati, Colin Mathers, Juan Rivera, Maternal, Child Undernutrition Study Group, et al. Maternal and child undernutrition: global and regional exposures and health consequences. *The lancet*, 371(9608):243–260, 2008.

[Blacklaws, 2018] Christina Blacklaws. Algorithms: transparency and accountability. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 376(2128):20170351, 2018.

[C *et al.*, 2022] Dhanamjayulu C, Nizhal U N, Praveen Kumar Reddy Maddikunta, Thippa Reddy Gadekallu, Celestine Iwendi, Chuliang Wei, and Qin Xin. Identification of malnutrition and prediction of bmi from facial images using real-time image processing and machine learning. *IET Image Processing*, 16(3):647–658, 2022.

[Dukhi, 2020] Natisha Dukhi. Global prevalence of malnutrition: evidence from literature. In *Malnutrition*. IntechOpen, 2020.

[Grantham-McGregor, 1995] Sally Grantham-McGregor. A review of studies of the effect of severe malnutrition on mental development. *The Journal of nutrition*, 125:2233S–2238S, 1995.

[He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

[Heemann *et al.*, 2021] Markus Heemann, Rockli Kim, Sebastian Vollmer, and SV Subramanian. Assessment of undernutrition among children in 55 low-and middle-income countries using dietary and anthropometric measures. *JAMA Network Open*, 4(8):e2120627–e2120627, 2021.

[Huang *et al.*, 2017] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.

[Kendall *et al.*, 2018] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7482–7491, 2018.

[Khan *et al.*, 2022] Misaal Khan, Mahapara Khurshid, Mayank Vatsa, Richa Singh, Mona Duggal, and Kuldeep Singh. On ai approaches for promoting maternal and neonatal health in low resource settings: A review. *Frontiers in Public Health*, 10:1864, 2022.

[Kuczmarski, 2002] Robert J Kuczmarski. *2000 CDC Growth Charts for the United States: methods and development*. Number 246. Department of Health and Human Services, Centers for Disease Control and Prevention, National Center for Health Statistics, 2002.

[Lie *et al.*, 2004] RK Lie, E Emanuel, C Grady, and D Wendler. The helsinki declaration, 2000, and ethics of human research in developing countries. *Indian Journal of Medical Ethics*, 1(3), 2004.

[Mathur *et al.*, 2018] Roli Mathur, Soumya Swaminathan, et al. National ethical guidelines for biomedical & health research involving human participants, 2017: A commentary. *The Indian journal of medical research*, 148(3):279, 2018.

[Ogden *et al.*, 2002] Cynthia L Ogden, Robert J Kuczmarski, Katherine M Flegal, Zuguo Mei, Shumei Guo, Rong Wei, Laurence M Grummer-Strawn, Lester R Curtin, Alex F Roche, and Clifford L Johnson. Centers for disease control and prevention 2000 growth charts for the united states: improvements to the 1977 national center for health statistics version. *Pediatrics*, 109(1):45–60, 2002.

[Riddick, 2003] Frank A Riddick. The code of medical ethics of the american medical association. *Ochsner Journal*, 5(2):6–10, 2003.

[Sethi *et al.*, 2020] Ashish Sethi, Mayank Vatsa, and Richa Singh. *Detecting Malnourishment in Children: Small Sample Size Problem in Deep Learning World*. Mtech thesis, IIIT-D-MTech-CS-AI-20-MT18024, 2020.

[Shrimpton *et al.*, 2001] Roger Shrimpton, Cesar G Victora, Mercedes de Onis, Rosangela Costa Lima, Monika Blossner, and Graeme Clugston. Worldwide timing of growth faltering: implications for nutritional interventions. *Pediatrics*, 107(5):e75–e75, 2001.

[Simonyan and Zisserman, 2015] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR,*, 2015.

[Wang *et al.*, 2020] Yaqing Wang, Quanming Yao, James T Kwok, and Lionel M Ni. Generalizing from a few examples: A survey on few-shot learning. *ACM computing surveys (csur)*, 53(3):1–34, 2020.

[WHO, 1995] WHO. *Physical status: The use of and interpretation of anthropometry, Report of a WHO Expert Committee*. World Health Organization, 1995.

[WHO, 2021] WHO. Ethics and governance of artificial intelligence for health: Who guidance. 2021.