DomainAdapt: Leveraging Multitask Learning and Domain Insights for Children's Nutritional Status Assessment

Misaal Khan[1,2][0009−0008−8030−6494], Richa Singh[1][0000−0003−4060−4573], Mayank Vatsa[1][0000−0001−5952−2274], and Kuldeep Singh[2]

[1] Indian Institute of Technology Jodhpur, India
[2] All India Institute of Medical Science Jodhpur, India
{khan.9, mvatsa, richa}@iitj.ac.in, kulpra@gmail.com

Abstract. This study presents a novel approach for automating nutritional status assessments in children, designed to assist health workers in public health contexts. We introduce "*DomainAdapt*," a novel dynamic task-weighing method within a multitask learning framework, which leverages domain knowledge and Mutual Information to balance task-specific losses, enhancing the learning efficiency for nutritional status screening. We have also assembled an unprecedented dataset comprising 16,938 multipose images and anthropometric data from 2,141 children across various settings, marking a significant first in this domain. Through rigorous testing, this method demonstrates superior performance in identifying malnutrition in children and predicting their anthropometric measures compared to existing multitask learning approaches. Dataset is available at : iab-rubric.org/resources/healthcare-datasets/anthrovision-dataset

Keywords: Malnourishment · Multi-task Loss Balancing · Anthropometry · Public Health.

1 Introduction

In the context of malnutrition detection, traditional methods include measuring height, weight, gender, Mid-Upper Arm Circumference (MUAC), and age (anthropometric measurements), recorded by trained healthcare professionals. These measurements are utilized to assess nutritional status indicators such as stunting, wasting, and underweight conditions based on weight-for-age (WFA), height-for-age (HFA), and BMI-for-age (BMI-forage) Z-scores. Following the World Health Organization (WHO) and Centers for Disease Control and Prevention(CDC) guidelines and growth charts, these assessments categorize each child as 'Healthy' or 'Unhealthy' [23,11]. However, these screening methods pose challenges in terms of the subjectivity of the worker, the time-consuming process of measuring each aspect one by one and the lack of scalability [16,17,4,1]. Researchers have explored machine learning (ML) based techniques on the point

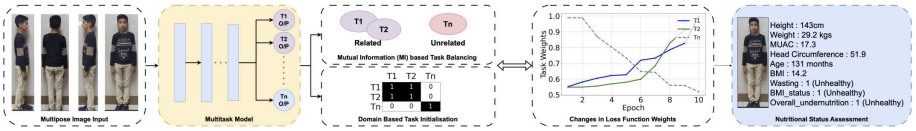

Fig. 1: Visual Representation of the developed Nutritional Status Screening Framework using DomainAdapt training. Input: Multipose images of the child, fed into a multitask model. Output: a detailed nutritional status assessment, including height, weight, MUAC, head circumference, age, BMI, and classifications for wasting, BMI status, and overall undernutrition

after recording the measurements that offer improved child malnutrition screening [12,13,20,6,15,7], however, they still suffer from the long and non-trivial requirement of recording each measurement one by one.

With the emergence of computer vision (CV) and artificial intelligence (AI) technologies, it is now possible to generate the aforementioned measurements from an image of a person [5,22]. Nonetheless, in such traditional techniques, issues like handling variations in camera distance, human orientation, and body posture often lead to inaccuracies in estimating the anthropometric measurements. These methods rely on pixel-based analysis, which struggles to account for the 3-D nature of human bodies and their diverse poses. Additionally, integrating multiple features such as face, body, and deep features into a cohesive model is challenging and computationally intensive [10,2,8]. To address these limitations, Multitask Learning (MTL) offers a more robust solution by simultaneously training the model on related tasks, thereby improving generalization and efficiency. In MTL, dynamic task weighting, such as uncertainty-based and gradient normalisation based methods, optimizes task influence and learning efficiency, reflecting a shift towards adaptive, responsive MTL systems [19,21]. Therefore, MTL allows for shared learning across tasks, enhancing the model's ability to handle variability in input data and improving overall performance in estimating the anthropometric measurements from images. However, research is hampered by the lack of comprehensive datasets, particularly for multi-pose image analysis, which is critical for reducing misclassification [25,18].

To address these challenges, this study aims to explore a new avenue of MTL based on shared information between the layers in enabling automated nutritional status screening for children based on their images, as shown in Figure 1. We test the proposed MTL approach by plugging it in a custom made dual-branched multitask model that employs VGG-19 for regression based anthropometric measurement extraction and ResNet-18 for classification. We highlight the use of domain knowledge and mutual information to enhance automation by increasing model efficiency and scalability despite the scarcity of literature and datasets. Hence, the developed framework does not require the health worker to have prior training aiding the under-resourced and remote areas. The contributions can be summarised as follows.

- Introduced "*DomainAdapt*", a dynamic MTL based training technique that optimizes task balance by integrating domain knowledge and Mutual Information, available open-source.
- Developed a DomainAdapt-based model for testing and evaluating the proposed MTL based technique for nutritional status screening from children's multi-pose images.
- Compiled a comprehensive dataset, named "*AnthroVision*" with diverse multi-pose child images and detailed anthropometrics for public use. The dataset is collected with multiple settings showcasing diverse backgrounds, clothing, and lighting conditions, comprising 2,142 subjects (1,326 in clinical and 815 in community settings). The clinical setting is captured at a hospital named All India Institute of Medical Sciences Jodhpur (AIIMS-J), India, and the community setting is captured in multiple government schools of the district.
- Evaluated *DomainAdapt*'s in a comparative analysis against traditional task weighting methods [3,9], showcasing improved learning outcomes and class imbalance handling.

2 Methodology

2.1 *AnthroVision*: Details of Dataset

To address the scarcity of publicly available datasets for automated nutritional status assessment in children, we introduce a novel dataset, "*AnthroVision*". Let's denote it $\mathcal{D}$. This dataset is a comprehensive collection of data points i, each represented by a tuple (I_i, yr_i, yc_i), where:

- I_i is the set of multi-angle images for the i-th subject, encompassing both full-body and facial views.
- yr_i is the vector of anthropometric measurements for the i-th subject, including age, height, weight, Body Mass Index (BMI), Mid-Upper Arm Circumference (MUAC), and Head Circumference (HC), serving as regression labels. Additionally, we have calculated the weight-for-age, height-for-age, and BMI-for-age z-scores based on the World Health Organization (WHO) [23] and the Centers for Disease Control and Prevention (CDC) [14] growth standards which are used for creating classification labels.
- yc_i is classified into categories such as healthy/unhealthy based on BMIzs-core, underweight or wasting (inferred from agewise MUAC cutoffs and weight-for-age < -2 SD), overall category (based on presence of stunting i.e. height-for-age Z-score (HAZ) < -2SD or underweight or wasting) [24]. Adhering to WHO and CDC standards ensures our dataset's validity and alignment with clinical practices in nutritional status assessment.

"*AnthroVision*", the proposed dataset is illustrated in Fig. 2 with settings showcasing diverse backgrounds, clothing, and lighting conditions, comprises 2,142 subjects (1,326 in clinical and 815 in community settings). The clinical setting is captured at a hospital named All India Institute of Medical Sciences

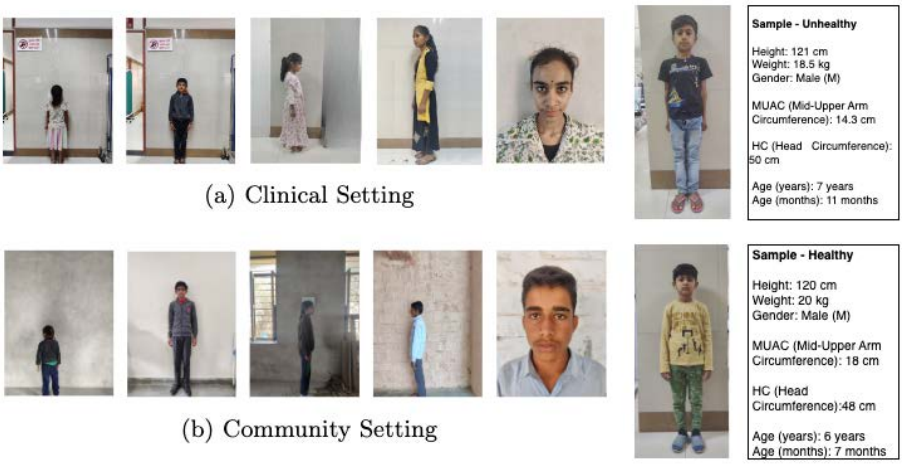

Fig. 2: *AnthroVision* Dataset: (a) Clinical Setting (Hospital - AIIMS Jodhpur, India)(b) Community Setting: nearby government schools in Jodhpur, Rajasthan, India district.

Jodhpur (AIIMS-J), India, and the community setting is captured in multiple government schools of the district. It is divided into 1284 training, 429 validation, and 428 test samples following a 60-20-20 split, with each subject contributing approximately eight multi-pose images. *AnthroVision*'s development involved utilizing standardized equipment and protocols in collaboration with government institutions, hospitals, and schools, with institutional ethical clearance from both our institute and the institute of data collection (wherever applicable). Further methodological details are provided in the supplementary materials.

2.2 *DomainAdapt*: Domain Guided Task Weighing Technique

In this section, the working of the proposed training technique, *DomainAdapt*, is explained. It is to be noted that when dealing with multitask learning, the interplay between tasks plays a pivotal role in model performance. Recognizing this, we introduce a novel learning technique, **DomainAdapt**, that incorporates domain-guided task weighting to adjust the importance of each task during training dynamically. This method leverages predefined domain knowledge and data-derived insights in the form of MI between predicted outcomes to optimize task-relatedness and performance.

Core Idea: When multiple poses of the user are taken via images, the images are used to predict the anthropometric values, which will be referred to as follows: Age (months) as T1, BMI as T2, Height(cm) as T3, Weight(kg) as T4, MUAC(cm) as T5, HC(cm) as T6. Further, image-based classification is done for health status based on nutritional indices BMIzscore-based Healthy/Unhealthy

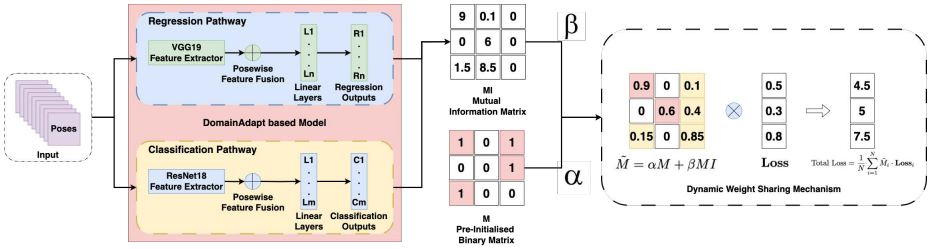

Fig. 3: Overview of the ***DomainAdapt*** Model with Regression and Classification pathways. Pose images are processed by VGG19 and ResNet18 feature extractors in two parallel pathways. The outputs from both pathways compute the mutual information (MI) matrix. A predefined binary matrix M and the MI matrix are combined using learnable parameters α and β, forming a dynamic weight matrix $\tilde{M} = \alpha M + \beta MI$ which weighs multitask losses

classification as T7, Underweight/Wasting (Based on agewise MUAC) classification as T8, and Overall (based on presence of Stunting/Underweight/Wasting) classification as T9 (0: healthy, 1: unhealthy).

***DomainAdapt* based Model:** We introduce a dual-branched multitask framework, as shown in Figure 3, that employs VGG-19 for predicting anthropometric measurements (Tasks T1-6) and ResNet-18 for classifying health status label (Tasks T7-9) (labels calculated based on nutritional indices). The architecture comprises two parallel feature extraction backbones: **VGG19** and **ResNet18**. We further perform the pose wise feature fusion for the task outputs of each pose and use this fused vector in the subsequent linear prediction layers. **Mutual Information Calculation and Dynamic Weighting:** To implement dynamic loss weighting, we utilize a binary matrix M of size 9×9, predefined based on domain knowledge, to represent task relationships. A value of '1' indicates related tasks, while '0' indicates unrelated tasks. For instance, in the context of BMI, relationships are established based on its dependency on Height and Weight. Thus, the matrix M is initialized with a value of '1' for Height and Weight and '0' for other metrics. This matrix serves as the foundation for dynamic weight adjustments during training, leveraging Mutual Information (MI) for optimization.

The mutual information matrix MI of size 9×9 is calculated for all task pairs. First, regression outputs are discretized into bins. Let $\mathbf{y} \in \mathbb{R}^n$ be the continuous regression outputs, and let B be the number of bins. The discretization function $\mathcal{D}$ maps the continuous outputs to discrete values:

$$\mathbf{y}_{\text{disc}} = \mathcal{D}(\mathbf{y}, B) \tag{1}$$

where $\mathbf{y}_{\text{disc}} \in \{1, 2, \ldots, B\}^n$, and $\mathcal{D}$ is defined as:

$$\mathcal{D}(\mathbf{y}, B) = \left\lfloor \frac{\mathbf{y} - \min(\mathbf{y})}{\max(\mathbf{y}) - \min(\mathbf{y})} \times B \right\rfloor \tag{2}$$

The classification outputs, which were originally probabilities, are then converted to binary values. Furthermore, the mutual information between tasks i and j is calculated as:

$$MI(i,j) = \sum_{x \in X} \sum_{y \in Y} p(x,y) \log \frac{p(x,y)}{p(x)p(y)} \tag{3}$$

where X and Y are the sets of discrete values of the outputs for tasks i and j, respectively, $p(x,y)$ is the joint probability distribution, and $p(x)$ and $p(y)$ are the marginal probability distributions.

Learnable Parameters α and β: These parameters are learned during the training process to optimize the performance of the model. (i) α: This parameter scales the contribution of the predefined binary matrix M. It reflects the initial assumptions and domain knowledge about the task relationships. (ii) β: This parameter scales the contribution of the mutual information matrix MI. It reflects the data-driven relationships discovered during training.

The dynamic weight matrix $\tilde{M}$ is calculated using these learnable parameters, and to ensure stable weight updates, $\tilde{M}$ is normalized:

$$\tilde{M} = \alpha M + \beta MI \quad \text{and} \quad \tilde{M}_{\text{norm}} = \frac{\tilde{M}}{\sum_{i,j} \tilde{M}(i,j)} \tag{4}$$

Task-Specific Loss Calculation: Task-specific losses, organized in a 9×1 vector, utilize Mean Squared Error (MSE) for regression tasks and Binary Cross-Entropy (BCE) for classification tasks. These losses are weighted by $\tilde{M}_{\text{norm}}$. The total loss for multitask learning is then calculated as:

$$\text{Total Loss} = \frac{1}{N} \sum_{i=1}^{N} \tilde{M}_{\text{norm},i} \cdot \mathbf{Loss}_i \tag{5}$$

where $\mathbf{Loss}_i$ denotes the individual loss for task i, and N is the total number of tasks.

3 Experiments and Results

Key hyperparameters include a batch size of 128 for training, 64 for validation, and 64 for testing, with an initial learning rate of 1e-3 using the Adam optimizer. Data augmentation techniques include color jitter, and random grayscale, which will not impact the shape of the human body. The model is trained on a single Tesla V100 GPU with CUDA 12.0. A 4-fold stratified cross-validation is applied.

Performance Evaluation. A comprehensive performance evaluation of the suggested training methodology for the said problem statement of predicting anthropometric measurements and nutritional labels is shown in Table 1. We have used *Precision*, *Recall* and *F1 score* to demonstrate the *classification* efficiency of the comparisons and *Root Mean Square Error(RMSE)* for assessing *regression*

Table 1: Performance Evaluation of *DomainAdapt* Model and Comparisons with Baseline Methods: Tasks T1-T6 are regression tasks representing Age (months), BMI, Height (cm), Weight (kg), MUAC (cm), and HC (cm), respectively. Tasks T7-T9 are classification tasks: T7 is BMIzscore-based Healthy/Unhealthy classification, T8 is Underweight/Wasting classification, and T9 is Overall classification based on stunting, underweight, and wasting. Precision, recall, and F1 score are used for classification tasks, while RMSE is used for regression tasks.

A. DomainAdapt performance comparison with other task weighing techniques for classification (Precision ↑, Recall ↑, F1 ↑) tasks

Model	T7			T8			T9		
	Precision ↑	Recall ↑	F1 ↑	Precision ↑	Recall ↑	F1 ↑	Precision ↑	Recall ↑	F1 ↑
DomainAdapt	0.51 ± 0.02	0.62 ± 0.01	0.54 ± 0.02	0.62 ± 0.05	0.69 ± 0.10	0.64 ± 0.14	**0.63 ± 0.01**	**0.67 ± 0.04**	**0.64 ± 0.03**
Static	**0.59 ± 0.11**	0.62 ± 0.07	**0.60 ± 0.15**	**0.67 ± 0.08**	**0.72 ± 0.03**	**0.67 ± 0.11**	0.62 ± 0.20	0.62 ± 0.14	0.62 ± 0.21
Uncertainty	0.54 ± 0.03	**0.67 ± 0.02**	0.56 ± 0.04	0.54 ± 0.05	0.73 ± 0.04	0.62 ± 0.05	0.49 ± 0.04	0.70 ± 0.03	0.58 ± 0.02
GradNorm	0.58 ± 0.03	0.54 ± 0.02	0.55 ± 0.03	0.53 ± 0.04	0.72 ± 0.03	0.61 ± 0.04	0.60 ± 0.03	0.61 ± 0.02	0.61 ± 0.03

B. Cross validation of model across clinical and community settings for classification (Precision ↑, Recall ↑, F1 ↑) tasks

Clinical	0.41 ± 0.02	0.64 ± 0.01	0.50 ± 0.03	0.56 ± 0.04	0.62 ± 0.05	0.58 ± 0.02	0.54 ± 0.03	0.58 ± 0.01	0.55 ± 0.02
Community	**0.70 ± 0.03**	**0.75 ± 0.02**	**0.70 ± 0.04**	**0.83 ± 0.05**	**0.85 ± 0.02**	**0.83 ± 0.03**	**0.77 ± 0.04**	**0.80 ± 0.03**	**0.77 ± 0.02**

C. Model performance comparison with other task weighing techniques for regression (RMSE↓) tasks

Model	RMSE T1 ↓	RMSE T2 ↓	RMSE T3 ↓	RMSE T4 ↓	RMSE T5 ↓	RMSE T6 ↓
DomainAdapt	**43.76 ± 1.07**	**2.96 ± 0.14**	**22.02 ± 0.66**	**12.44 ± 0.22**	**3.55 ± 0.24**	**5.05 ± 1.07**
Static	44.02 ± 1.12	3.47 ± 0.21	28.87 ± 0.66	**12.31 ± 0.34**	4.32 ± 0.11	9.42 ± 1.25
Uncertainty	45.00 ± 1.23	3.86 ± 0.27	30.48 ± 1.02	12.74 ± 0.52	4.08 ± 0.25	10.36 ± 0.67
GradNorm	52.00 ± 1.34	4.58 ± 0.35	23.47 ± 1.12	27.06 ± 0.74	4.07 ± 0.31	28.51 ± 0.84

D. Cross validation of model across clinical and community settings for regression (RMSE↓) tasks

Clinical	45.02 ± 1.23	4.41 ± 0.31	29.55 ± 1.02	**12.30 ± 0.52**	5.25 ± 0.25	13.46 ± 0.67
Community	46.96 ± 1.34	5.47 ± 0.37	36.42 ± 1.56	17.20 ± 1.72	6.94 ± 1.31	15.23 ± 0.84

Table 1 **A** and **C** showcase *DomainAdapt* based model performance compared to other baseline MTL training techniques such as *Static*, *Uncertainty* and *Gradient Normalisation (GradNorm)*. DomainAdapt has shown somewhat at par performance at classification for Task 1 and 2 while showed superior classification for Task3. However, the important observation here lies in the regression performance where *DomainAdapt* has shown the best performance which underscores our hypothesis that DomainAdapt has the capability to improve overall multitask performance by balancing out related tasks and improving them holistically. Table 1 **B** and **D** shows the cross validation of the model's performance across clinical and community settings in contrast to a model that is trained on both the datasets combined. Interestingly, the model trained on Community setting has shown much superior performance in comparison to the other two while in case of regression, the model trained on a diverse dataset shows the least *RMSE*. This might indicate on the role of the kind of lighting, capturing environment and clothing on the result outcomes. Most subjects in a community setting are wearing a uniform as we have collected those datasets from various schools in the district and have varied backgrounds pertaining to the locations. The images collected in clinical environment are diverse in terms of clothing with consistency in the image background. Further ablation analysis of the impact of removing particular poses on model performance is explored in the *supplementary material*.

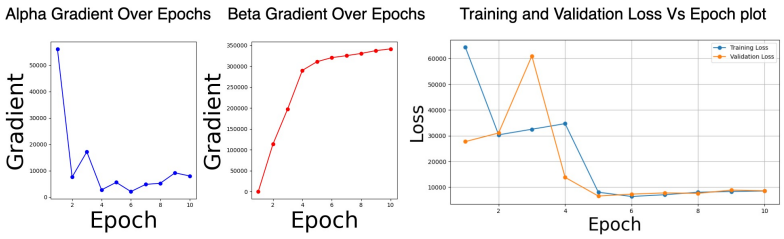

Fig. 4: Comparison of Alpha and Beta Gradients with Train and Validation Loss. (1) **Alpha Gradient Over Epochs** shows the gradient variation of α during training impacting M matrix. (2) **Beta Gradient Over Epochs** illustrates the gradient changes of β impacting MI matrix. (3) **Training and Validation Loss vs. Epoch Plot** compares training and validation loss over epochs to assess model performance and generalization.

3.1 Analysis of *DomainAdapt*

Plots of gradients for parameters α and β in Fig 4 show diverging trends, suggesting a shift from relying on the initial task importance matrix M to valuing the mutual information matrix MI more. This is confirmed by training and validation loss curves, where α's stabilization and β's increase hint at improved model generalization through a better balance of tasks and enhanced focus on dynamic task relationships.The range of α value is between 0.8-1 while for β it is 0.4-1. We examine our *DomainAdapt* framework by visualizing the task weight matrices from the multitask learning process in 3D graphs (in Fig. 5), showcasing weight distribution for nine tasks. Each graph's z-axis shows weight magnitude, with the x and y-axes detailing task-specific 'bins/attributes,' revealing the algorithm-driven variability in task importance.

The visualizations reveal a varied weight distribution across tasks, indicating the model's adaptive approach to prioritizing different tasks based on data. Patterns of unique weighting per task suggest tailored contributions to the learning goal. Notably, areas with high peaks could risk overfitting, while low-weight regions might underfit. This demonstrates the *DomainAdapt* framework's ability to dynamically balance task focus, promoting a well-rounded learning state that enhances generalization.

4 Conclusion

Our study advances and establishes state-of-the-art for the field of computer-assisted interventions for automating malnutrition screening in children to address global health inequities. Through the *DomainAdapt* multitask learning framework, we blended domain knowledge with computer vision and created a set of a unique, large-scale dataset. This highlights our commitment to developing accessible medical health solutions, particularly for vulnerable and underrepresented populations in remote areas where conventional screening techniques

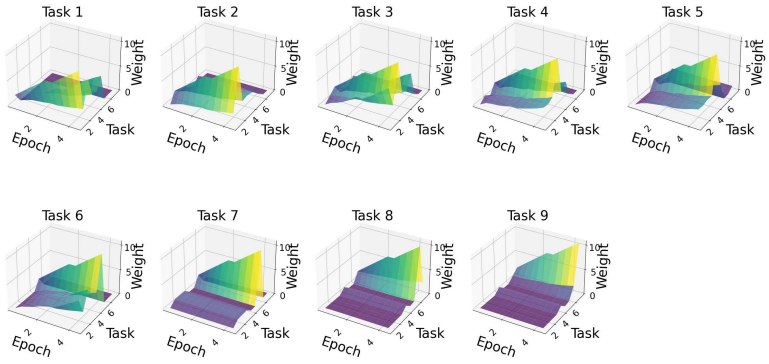

Fig. 5: Adaptive Weight Distributions Across Multitask Learning Objectives. Each plot (Task 1 to Task 9) illustrates how the weight distribution changes dynamically during the training process, reflecting the model's adaptation to varying task difficulties and interdependencies. X-axis: training epochs, Y-axis: task indices, and Z-axis: weight values.

are not very scalable due to involved costs of training healthworkers or the time taken in spanning a large geographical area. This research also paves the way for future innovations in point-of-care and teleradiology applications for complex problem statements without prior literature, promising a significant impact on public health strategies worldwide.

Acknowledgments. M.Khan is partly supported through Prime Minister's Research Fellowship and M.Vatsa is partially supported through SwarnaJayanti Fellowship by the Government of India.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Barker, L.A., Gout, B.S., Crowe, T.C.: Hospital malnutrition: prevalence, identification and impact on patients and the healthcare system. International journal of environmental research and public health **8**(2), 514–527 (2011)
2. Caruana, R.: Multitask learning. Machine learning **28**(1), 41–75 (1997)
3. Chen, Z., Badrinarayanan, V., Lee, C.Y., Rabinovich, A.: Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In: Proceedings of the 35th International Conference on Machine Learning. vol. 80, pp. 794–803. PMLR (2018)
4. Cross, T.L.: Social emotional needs: The effects of educational malnourishment on the psychological well-being of gifted students. Gifted Child Today **37**(4), 264–265 (2014)

5. Hassan, B., Sherazi, H.H.R., Ali, M., Siddiqi, Y.: Estimating anthropometric soft biometrics: An empirical method. Intelligent Automation & Soft Computing **37**(3) (2023)
6. How, M.L., Chan, Y.J.: Artificial intelligence-enabled predictive insights for ameliorating global malnutrition: A human-centric ai-thinking approach. AI **1**(1), 68–91 (2020)
7. Hussain, Z., Borah, M.D.: Nutritional status prediction in neonate using machine learning techniques: a comparative study. International Conference on Machine Learning, Image Processing, Network Security and Data Sciences pp. 69–83 (2020)
8. Kendall, A., Gal, Y., Cipolla, R.: 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. pp. 7482–7491 (2018)
9. Kendall, A., Gal, Y., Cipolla, R.: 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 (CVPR) (2018)
10. Khan, M., Agarwal, S., Vatsa, M., Singh, R., Singh, K.: Nutriai: Ai-powered child malnutrition assessment in low-resource environments. In: Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence. pp. 6378–6385 (2023)
11. Khan, M., Khurshid, M., Vatsa, M., Singh, R., Duggal, M., Singh, K.: On ai approaches for promoting maternal and neonatal health in low resource settings: A review. Frontiers in Public Health **10**, 880034 (2022)
12. Khare, S., Kavyashree, S., Gupta, D., Jyotishi, A.: Investigation of nutritional status of children based on machine learning techniques using indian demographic and health survey data. Procedia computer science **115**, 338–349 (2017)
13. Kraamwinkel, N., Ekbrand, H., Davia, S., Daoud, A.: The influence of maternal agency on severe child undernutrition in conflict-ridden nigeria: Modeling heterogeneous treatment effects with machine learning. PloS one **14**(1), e0208937 (2019)
14. Kuczmarski, R.J.: 2000 CDC Growth Charts for the United States: methods and development. No. 246, Department of Health and Human Services, Centers for Disease Control and Prevention, National Center for Health Statistics (2002)
15. Momand, Z., Zarinkhail, M.S., Aryan, M.F.: Machine learning based prediction of edematous malnutrition in afghan children. International Conference on Emerging Technologies and Intelligent Systems pp. 235–245 (2021)
16. Nahar, B., Hamadani, J.D., Ahmed, T., Tofail, F., Rahman, A., Huda, S., Grantham-McGregor, S.: Effects of psychosocial stimulation on growth and development of severely malnourished children in a nutrition unit in bangladesh. European Journal of Clinical Nutrition **63**(6), 725–731 (2009)
17. Reilly Jr, J.J., Hull, S.F., Albert, N., Waller, A., Bringardener, S.: Economic impact of malnutrition: a model system for hospitalized patients. Journal of Parenteral and Enteral nutrition **12**(4), 371–376 (1988)
18. Ruder, S., Bingel, J., Augenstein, I., Søgaard, A.: Latent multi-task architecture learning. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33, pp. 4822–4829 (2019)
19. Sener, O., Koltun, V.: Multi-task learning as multi-objective optimization. In: Advances in Neural Information Processing Systems (2018)
20. Sharma, V., Sharma, V., Khan, A., Wassmer, D.J., Schoenholtz, M.D., Hontecillas, R., Bassaganya-Riera, J., Zand, R., Abedi, V.: Malnutrition, health and the role of machine learning in clinical setting. Frontiers in nutrition **7**, 44 (2020)

21. Standley, T., Zamir, A.R., Chen, D., Savarese, S., Malik, J., Savva, M.: Which tasks should be learned together in multi-task learning? In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6428–6437 (2020)
22. Sui, J., Bu, C., Zhao, X., Liu, C., Ren, L., Qian, Z.: Body weight estimation using virtual anthropometric measurements from a single image. IEEE Transactions on Instrumentation and Measurement (2023)
23. WHO: The who child growth standards, https://www.who.int/tools/child-growth-standards/standards
24. World Health Organization: Malnutrition in children. https://www.who.int/data/nutrition/nlis/info/malnutrition-in-children (2024), accessed: 2024-03-01
25. Yang, Y., Hospedales, T.M.: Deep multi-task learning with shared memory. In: European Conference on Computer Vision. pp. 605–621. Springer (2016)