

AI-Based Expanded BMI System for Early Malnutrition Risk Detection

S. V. Rajeswari¹, Namitha M², Monika P³, Nagadharshini A⁴, Usha P⁵

¹ M.E., Assistant Professor, Department of CSE, Mahendra Engineering College for Women, Tiruchengode, Tamil Nadu, India

^{2,3,4,5} Department of Computer Science and Engineering, Mahendra Engineering College for Women, Tiruchengode, Tamil Nadu, India

Abstract - Childhood malnutrition continues to be a pressing concern in public health, especially in regions where trained nutrition personnel and diagnostic resources remain limited. This paper presents an automated image-based screening system that employs the YOLOv5m deep learning algorithm to identify malnutrition in children from a single photograph and subsequently delivers age-appropriate dietary guidance without requiring specialist intervention. The detection model was trained on 3,200 annotated paediatric images sourced from Kaggle, applying transfer learning initialised from MS COCO weights combined with a mosaic augmentation and jitter pipeline. Evaluated on a 320-image held-out test partition, the system attains 94.7% accuracy, 93.2% precision, 95.1% recall, and an AUC-ROC of 0.974, outperforming ResNet-50 and VGG-16 baselines by 5.1 and 7.5 percentage points respectively. A rule-driven recommendation engine aligned with WHO IMAM and ICMR-RDA 2020 guidelines translates each detected nutritional deficit into ranked meal plans filtered by age group and allergen profile. Deployed through a Flask REST backend and an HTML/CSS/JavaScript front-end, the system achieves a mean inference time of 187 ms per image on commodity hardware, rendering it practical for resource-limited primary health settings.

Key Words: YOLO, childhood malnutrition, deep learning, paediatric screening, nutritional recommendation, image classification, transfer learning.

I. INTRODUCTION

Malnutrition among children under the age of five remains one of the foremost preventable health burdens worldwide. The World Health Organization reports that over 230 million children globally are affected by stunting, acute wasting, or micronutrient deficiencies, with the greatest concentration in South Asia and sub-Saharan Africa [15]. The downstream effects of undetected malnutrition include irreversible neurocognitive impairment, poor school performance, and reduced adult productivity, collectively imposing an economic cost estimated at two to three percent of gross domestic product annually on high-burden nations.

Conventional screening depends on anthropometric measurements such as weight-for-height z-score, height-for-age z-score, and mid-upper arm circumference. Each procedure demands a trained operator, calibrated instruments, and a cooperative subject. Published studies document misclassification rates as high as 15 percent when measurements are performed by community volunteers rather than qualified clinical staff. Computer vision offers a technically viable path: visual biomarkers of malnutrition are perceptible in standard photographs, and deep convolutional networks have demonstrated near-specialist sensitivity for analogous visual recognition tasks in dermatology and diabetic retinopathy screening [22].

II. LITERATURE SURVEY

A. Deep Learning for Malnutrition and Growth Assessment

Lakshminarayanan et al. [1] investigated convolutional neural networks for automated malnutrition identification, applying AlexNet to a paediatric photograph dataset and reporting classification accuracy exceeding 90 percent. Ankalaki et al. [2] assembled a four-class annotated dataset and evaluated multiple deep learning architectures to differentiate severity grades, finding that mosaic-style data augmentation substantially reduced overfitting. Biradar and Naik [3] applied ResNet-18 to paediatric nutritional status classification, achieving 95 percent accuracy, demonstrating that transfer learning from ImageNet weights converged faster than random initialisation.

B. YOLO-Based Detection in Healthcare

Redmon et al. [4] introduced the single-pass YOLO detection paradigm enabling real-time throughput impractical with region-proposal methods. Jocher et al. [5] released YOLOv5 with a CSPDarknet53 backbone and PANet neck; the medium variant YOLOv5m balances detection capability with computational overhead suitable for modest GPU resources. Diwan et al. [6] noted that healthcare deployments particularly benefit from YOLO's real-time inference and open-source pre-trained weights which facilitate domain-specific fine-tuning where labelled data are scarce.

C. Nutritional Recommendation Systems

Toledo et al. [7] developed a food recommendation platform integrating nutritional databases with user dietary preferences, demonstrating measurable improvements in dietary compliance when personalised recommendations were delivered via a web interface. Mansouri et al. [8] concluded that end-to-end pipelines coupling visual classification with structured nutritional knowledge bases produce more actionable outputs than standalone detection models.

III. EXISTING SYSTEM

Current field-level practice for paediatric nutritional assessment relies on manual anthropometric procedures unchanged for several decades. Frontline healthcare workers record weight, height, and mid-upper arm circumference, then compare against WHO 2006 Child Growth Standard reference tables. This workflow has structural weaknesses: uncooperative infants require both training and practice to measure accurately; calibrated equipment is expensive and frequently absent in peripheral facilities. Converting measurements into personalised dietary guidance requires printed lookup tables or a connected health information system, neither reliably available at community tier.

IV. DISADVANTAGES OF EXISTING SYSTEM

Dependence on trained clinical personnel restricts deployment to staffed health facilities, leaving large rural populations without routine nutritional surveillance. The multi-step measurement process delays identification of at-risk children. Instrument calibration errors and operator inconsistency introduce classification noise, with documented false-negative rates as high as 15 percent in volunteer-operated surveys. No existing field-level tool automatically translates a detected nutritional deficit into a personalised dietary prescription, creating a systematic gap between screening and actionable follow-up.

V. PROPOSED SYSTEM

The proposed system replaces manual anthropometric procedures with an automated, photograph-driven pipeline. A healthcare worker or parent captures a standard image of the child using any smartphone; the photograph is uploaded to the web interface and forwarded to the detection backend without requiring specialised equipment. A fine-tuned YOLOv5m model classifies the child as malnourished or healthy within a single forward pass. On a positive detection, the system queries a structured nutritional knowledge base and returns ranked age-appropriate meal recommendations aligned with ICMR-RDA 2020 guidelines and filtered by allergen profile. The Flask REST backend serves both model inference and database queries through a single API endpoint, compatible with low-bandwidth rural internet connections.

VI. SYSTEM REQUIREMENTS

Hardware Requirements

Processor: Intel Core i5 or equivalent (minimum 2.4 GHz). RAM: Minimum 1 GB (4 GB recommended). Storage: 160 GB secondary storage (includes model weights 14.3 MB and knowledge base). Camera: Any smartphone camera capable of capturing standard child photographs.

Software Requirements

OS: Windows 10 / Ubuntu 20.04 or later. Language: Python 3.9. Framework: Flask 2.3, PyTorch 2.0, Ultralytics YOLOv5. Database: SQLite via SQLAlchemy 2.0 ORM. Frontend: HTML5, CSS3, JavaScript (ES2020), Bootstrap 5.3. IDE: PyCharm / VS Code.

VII. SYSTEM ARCHITECTURE

The architecture comprises five sequentially coupled processing layers. (1) Data Acquisition Layer: child images annotated in YOLO format using LabelImg; inter-annotator Cohen's kappa of 0.87 confirmed labelling consistency. (2) Pre-processing Layer: all inputs resized to 640x640 pixels and normalised to [0, 1] using an 80:10:10 stratified split. (3) Detection Layer: YOLOv5m processes each tensor through CSPDarknet53 backbone, PANNet neck, and three-scale detection head; NMS with IoU threshold 0.45 consolidates overlapping predictions. (4) Recommendation Layer: SQLAlchemy ORM queries SQLite knowledge base by predicted class and age band, with allergen bitfield filter enabling O(1) retrieval. (5) Presentation Layer: Flask REST endpoint returns JSON consumed by the HTML/CSS/JavaScript front-end, updating the result panel without a full page reload.

VIII. METHODOLOGY

A. Dataset and Annotation

The Kaggle paediatric malnutrition dataset contains 3,200 photographs of children aged 0-60 months, balanced at 1,600 images per class, spanning diverse backgrounds, lighting conditions, subject poses, and camera distances. Bounding-box annotations were produced with LabelImg and independently reviewed by two annotators; conflicts resolved by a third reviewer. The consolidated annotation set attained a mean IoU of 0.91 relative to the consensus ground truth.

B. Augmentation Strategy

Standard horizontal flip ($p=0.5$) and brightness-contrast jitter ($\pm 20\%$) were supplemented with mosaic augmentation compositing four images into a single 640x640 frame. Gaussian noise injection ($\sigma=0.01$) addressed low-megapixel sensor noise. HSV colour-space jitter (hue ± 0.015 , saturation ± 0.7 , value ± 0.4) extended skin-tone representation across the demographic diversity of the dataset.

C. Model Training

YOLOv5m was trained using SGD with initial learning rate 0.01 cosine-annealed to 0.001 at epoch 100, momentum 0.937, weight decay 5×10^{-4} , batch size 16. Backbone initialised from MS COCO weights; neck and detection head fine-tuned over a 10-epoch warm-up before full-network release. Checkpoint at epoch 82 ($mAP@0.5 = 0.961$) was retained. Total training time on NVIDIA RTX 3060: 4.2 hours.

D. Inference Pipeline and Recommendation Engine

Uploaded images are padded to square aspect ratio, resized to 640x640, and normalised. Predictions with confidence below 0.50 are suppressed. The Flask endpoint returns class label, bounding-box coordinates, and confidence score as JSON within 187 ms. SQLAlchemy ORM queries use predicted class and age band as composite primary keys; top five items per meal slot scored by nutrient density relative to ICMR-RDA 2020 requirements are rendered as a formatted table.

IX. IMPLEMENTATION

The system was developed on Windows OS using Python 3.9 within PyCharm IDE. The back-end is implemented as a Flask 2.3 REST application; YOLOv5m inference is executed through the Ultralytics Python package with PyTorch 2.0. The relational knowledge base is managed via SQLite and accessed through SQLAlchemy 2.0 ORM without a separate database server. The front-end uses HTML5, CSS3, and vanilla JavaScript (ES2020) with the Fetch API and FormData; the JSON response updates the DOM dynamically without a page reload. The complete application including model weights (14.3 MB), knowledge base, and web assets runs on an Intel Core i5 at 2.6 GHz with 1 GB RAM.

X. RESULTS AND DISCUSSION

The system was evaluated on a held-out test set of 320 images (160 per class). The proposed YOLOv5m model achieves the highest scores across all five metrics. Its 5.1 percentage-point accuracy advantage over ResNet-50 and 7.5 points over VGG-16 confirm the single-pass YOLO architecture with mosaic augmentation and domain-specific fine-tuning is better suited to this detection task than standard classification backbones. The AUC-ROC of 0.974 indicates strong discriminative capability critical for a screening application where false negatives carry clinical consequences. Inference latency averages 187 ms, within the sub-200 ms

threshold required for a responsive web interface. The recommendation engine returns dietary suggestions in under 3 ms per query.

Table -1: Comparative Classification Performance (Test Set, n = 320)

Model	Acc. (%)	Prec. (%)	Rec. (%)	F1 (%)	AUC- ROC
VGG-16	87.2	86.4	88.1	87.2	0.923
ResNet-50	89.6	88.9	90.3	89.6	0.941
MobileNet V2	86.1	85.7	86.8	86.2	0.912
Proposed (YOLOv5m)	94.7	93.2	95.1	94.1	0.974

XI. ADVANTAGES OF THE PROPOSED SYSTEM

The proposed system eliminates dependence on calibrated instruments and trained measurement personnel, enabling malnutrition screening at any location with a camera-equipped device. Automated classification reduces human error and ensures consistent output regardless of operator experience. Real-time inference within 187 ms on commodity hardware supports deployment in settings with limited computational infrastructure. The integrated nutritional recommendation engine bridges the gap between screening and intervention by providing actionable dietary guidance immediately upon a positive detection.

XII. CONCLUSION

This paper has presented an automated, image-based malnutrition detection system applying a fine-tuned YOLOv5m deep learning model to classify children as malnourished or healthy from a single photograph. The system integrates detection, classification, and nutritional recommendation into a unified web application deployable on standard consumer hardware without specialist involvement. Experimental evaluation demonstrated 94.7% accuracy, 93.2% precision, 95.1% recall, and an AUC-ROC of 0.974, surpassing ResNet-50 and VGG-16 baselines by statistically meaningful margins. The system runs at 187 ms per image on an Intel Core i5 with 1 GB RAM, confirming practical feasibility for resource-limited primary health settings.

XIII. FUTURE WORK

Future development will prioritise dataset expansion to 15,000+ annotated images representing South Asian, sub-Saharan African, and Southeast Asian populations, enabling multi-class severity grading. The recommendation engine will be upgraded to a retrieval-augmented generation pipeline over a large language model. Offline inference on low-end Android devices via INT8 post-training quantisation through ONNX Runtime will be investigated. A prospective observational study in collaboration with the Tamil Nadu ICDS network is planned to validate sensitivity and specificity across 1,200 children.

REFERENCES

- [1] A. R. Lakshminarayanan et al., "Malnutrition detection using convolutional neural network," in Proc. 7th Int. Conf. Bio Signals, Images, and Instrumentation (ICBSII), IEEE, 2021, pp. 1–5.
- [2] S. Ankalaki, V. G. Biradar, K. K. Naik, and G. S. Hukkeri, "A deep learning approach for malnutrition detection," Int. J. Online Biomed. Eng. (iJOE), vol. 20, no. 6, pp. 116–138, 2024.
- [3] V. G. Biradar and K. K. Naik, "Detection of malnutrition in children using deep learning model," in Smart Trends in Computing and Communications (SmartCom 2024), vol. 946, Springer, 2024.
- [4] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, real-time object detection," in Proc. IEEE CVPR, Las Vegas, NV, Jun. 2016, pp. 779–788.

- [5] G. Jocher et al., "Ultralytics YOLOv5," Zenodo, 2021, doi: 10.5281/zenodo.3908559.
- [6] T. Diwan, G. Anirudh, and J. V. Tembhurne, "Object detection using YOLO: challenges, architectural successors, datasets and applications," Multimed. Tools Appl., vol. 82, no. 6, pp. 9243–9275, 2023.
- [7] R. Y. Toledo, A. A. Alzahrani, and L. Martinez, "A food recommendation system considering nutritional information and user preferences," IEEE Access, vol. 7, pp. 96695–96711, 2019.
- [8] M. Mansouri et al., "Deep learning for food image recognition and nutrition analysis," SN Comput. Sci., vol. 4, no. 513, 2023.
- [15] World Health Organization, "Malnutrition," WHO Fact Sheet, Geneva, Jun. 2023.
- [17] Indian Council of Medical Research, "Nutrient Requirements for Indians: RDA and EAR," ICMR-NIN, Hyderabad, 2020.
- [22] A. Esteva et al., "Dermatologist-level classification of skin cancer with deep neural networks," Nature, vol. 542, no. 7639, pp. 115–118, Feb. 2017.