International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 04 April 2026
Page No: 2165-2171
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0289

AI-Driven BMI Estimation and Personalised Health Recommendation System

Ms. V. Deepa Priya [1], Mr. M. H. Farseel [2], Mr. M. Xavier Raj [3]
[1]Assistant Professor, Information Technology, Kamaraj College of Engineering and Technology, Madurai, Tamilnadu, India, 625 701
[2,3]UG, Information Technology, Kamaraj College of Engineering and Technology, Madurai, Tamilnadu, India, 625 701
Email ID: *deepapriyait@kamarajengg.edu.in[1], 22uit048@kamarajengg.edu.in[2], 22uit124@kamarajengg.edu.in[3]*

Abstract
Access to basic health metrics such as Body Mass Index remains constrained for billions of people who lack proximity to clinical facilities. This paper presents an AI-driven system that estimates height, weight, and BMI from a single full-body photograph using two custom-trained U-Net neural networks operating under a multi-task learning framework on the IMDB-Wiki dataset. A pretrained Keypoint R-CNN model serves as an image quality gate before inference, generates personalised, BMI-category-specific health recommendations. The system is deployed across two production-ready platforms: a three-tier web application (React 18 + Express.js + FastAPI) and a native Android mobile application (React Native + Expo). Evaluated on 15 informal test subjects, the height model achieved a Mean Absolute Error of 6.136 cm and the weight model 9.8 kg, with 80% correct BMI category classification. All misclassifications occurred within three BMI units of a WHO category boundary. The system responds in under 3 seconds and demonstrates that meaningful preventive health screening can be delivered from a smartphone photograph alone, with no specialist equipment required.

Keywords: *BMI estimation; Deep learning; Health recommendation; Keypoint R-CNN; U-Net*

1. Introduction
Routine health screening sits at the foundation of preventive medicine, yet for a large proportion of the global population it remains out of reach. The World Health Organization's 2023 global health estimates indicate that over 1.9 billion adults are overweight and approximately 462 million remain underweight, yet fewer than one in five adults in low- and middle-income countries received a BMI assessment in the preceding year [1]. The barriers are both logistical — distance from facilities, cost, clinic hours — and cultural, with many people avoiding clinical settings until symptoms become acute. Body Mass Index, conceptualised by Adolphe Quetelet in 1832 and formalised as a WHO screening tool in 1995, remains the standard first-line proxy for body fat and health risk precisely because it requires no specialised equipment beyond a scale and stadiometer. Removing even those basic requirements through image-based estimation would extend the reach of health screening to populations that currently go unscreened. Recent advances in computer vision have demonstrated that photographs contain far more quantitative biometric information than casual observation reveals. Encoder-decoder architectures such as U-Net [2], originally developed for biomedical image segmentation, have been adapted for body measurement regression. Multi-task learning frameworks have been shown to regularise regression models and improve generalisation on biometric tasks [3]. Pose estimation models such as Keypoint R-CNN [4] enable automated quality control of input images, a practical requirement for real-world deployment. This paper presents a complete end-to-end integration of these components into a dual-platform health screening system that accepts a full-body photograph and returns height,

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 04 April 2026
Page No: 2165-2171
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0289

weight, BMI, and a personalised AI-generated health recommendation within seconds. The system is delivered as both a web application and a native Android application sharing a common machine learning backend, and is evaluated on test subjects with known measurements.

1.1.Background and Motivation

Despite the global scale of the BMI-related health challenge, regular screening remains inaccessible to large segments of the population. Early image-based measurement work relied on classical computer vision — silhouette extraction, edge detection, and statistical body shape models — which were brittle under real-world conditions. The advent of deep convolutional neural networks fundamentally changed the landscape. Cao et al. [5] introduced real-time multi-person pose estimation, and Gu et al. [3] proposed a multi-task CNN that simultaneously performed body segmentation and height regression, demonstrating that auxiliary tasks substantially improved regression accuracy. These findings

1.2.Method

The system was designed around a modular, dual-deployment architecture. A FastAPI machine learning service encapsulates all inference logic including pose validation, height and weight prediction, and BMI calculation. A React web application connects through an Express.js middleware layer that also manages the Gemini API key. A React Native Android application connects directly to FastAPI, reducing latency. This shared-backend design ensures prediction consistency across both platforms while minimising code duplication.

1.3.U-Net Architecture and Training

Both height and weight prediction models are based on U-Net — an encoder-decoder architecture with skip connections that preserve spatial detail through the network. The encoder comprises five ConvDown blocks, each applying two 3×3 convolutions with Batch Normalisation and ReLU, followed by 2×2 max pooling. The spatial resolution reduces from 256 to 8 pixels while channel depth grows from 3 to 2048. The decoder mirrors this through four ConvUpsample blocks, and the regression head collapses the output to a scalar via Flatten → Linear(32768→1024) + ReLU → Dropout(0.5) → Linear(1024→1). Both models were trained on the IMDB-Wiki dataset (~500,000 full-body photographs) under a multi-task learning framework, simultaneously optimising body segmentation (Dice loss), keypoint heatmap detection (cross-entropy), and biometric regression (MAE). The height model best checkpoint is model_ep_48.pth and the weight model checkpoint is model_ep_37.pth. directly motivated the architecture adopted in this work. Additionally, knowing one's BMI is only the first step; large language models such as Google Gemini are capable of generating contextually sensitive, personalised guidance far more actionable than static template text shown in Table 1 [6].

Table 1 U-Net Encoder Architecture — Channel

Block	Input Channels	Output Channels	Spatial Resolution
ConvDown 1	3 (RGB)	128	128 × 128 px
ConvDown 2	128	256	64 × 64 px
ConvDown 3	256	512	32 × 32 px
ConvDown 4	512	1024	16 × 16 px
ConvDown 5 (Bottleneck)	1024	2048	8 × 8 px

1.4. Pose Validation Using Keypoint R-CNN

Before any image reaches the U-Net models, it passes through a Keypoint R-CNN quality gate (ResNet-50 FPN backbone, pretrained on COCO Keypoints). The gate enforces three sequential checks: (1) detected person confidence > 0.40; (2) nose, both shoulders, and both hips visible; (3) both ankles visible in frame. Failure at any check returns an HTTP 400 response with a plain-English description, enabling users to retake the photograph. This quality gate was found to reduce real-world height prediction error by 3.6× compared to unfiltered inference.

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 04 April 2026
Page No: 2165-2171
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0289

1.5.Output Scaling and BMI Computation

The raw scalar output of each U-Net regression head is converted to physically meaningful units using empirically derived constants: height_cm = raw / −40.0 and weight_kg = raw / −275.0. Both values are validated against physiological bounds (height: 130–230 cm; weight: 40–150 kg) before BMI is computed as BMI = weight / (height in metres)², and a WHO category is assigned (Underweight < 18.5; Normal 18.5–24.9; Overweight 25.0–29.9; Obese ≥ 30.0).

1.6.Backend API

The FastAPI service loads all three models at startup and exposes a /predict endpoint (POST, multipart/form-data) that runs the full pipeline. The Express.js middleware handles file uploads via Multer (5 MB limit), assembles a Gemini prompt from the prediction output, and calls the Gemini API with a 10-second timeout, falling back to a category-specific local template on failure. Table 2 summarises the API endpoints.

Table 2 API Endpoint Reference

Service	Method	Endpoint	Description
FastAPI	POST	/predict	Pose validation + inference + BMI + fallback recommendation
FastAPI	GET	/	Service health check
Express.js	POST	/api/predict	Upload → FastAPI → Gemini → consolidated response
Express.js	GET	/api/health	Server status check

1.7.Frontend Implementations

The React 18 web application uses a three-component architecture (ImageUpload, PredictionResults, HealthRecommendations) rendered in a responsive TailwindCSS grid. A BMI scale bar implemented as a CSS linear gradient — blue through green, orange, and red, matching WHO category colours — gives users immediate visual context without needing to memorise threshold values. The React Native Android application (Expo SDK ~54, targeting Android API 34) manages all state in a single App.js using React hooks, selects images via expo-image-picker, and submits them to FastAPI via native Fetch as multipart/form-data. The Android APK was compiled using Expo's EAS cloud build service and side-loaded onto physical devices for testing.

2. Results And Discussion

2.1.Results

Mean Absolute Error (MAE) was used as the primary evaluation metric. The height model achieved a MAE of 6.136 cm and the weight model a MAE of 9.8 kg across the validation set. The estimated combined BMI MAE is approximately 3.5 units — sufficient for correct WHO category classification for subjects more than 5 BMI units from the nearest boundary. On a best-case test with a well-framed photograph (true height 174.7 cm, true weight 70.0 kg), the system predicted 174.02 cm and 69.85 kg, yielding a BMI error of 0.13 kg/m² with correct Normal weight classification. Table 3 summarises accuracy across all metrics.

Table 3. Model Prediction Accuracy — Validation MAE Results

Model	Metric	Value	Practical Interpretation
Height Model (model_ep_48.pth)	MAE	6.136 cm	Average prediction ~6.1 cm from true height
Weight Model (model_ep_37.pth)	MAE	9.8 kg	Average prediction ~9.8 kg from true weight
Combined (estimated)	BMI MAE	~3.5 units	Sufficient for category classification away from

			boundaries

Informal testing across 15 subjects with independently verified measurements showed that 73.3% of height predictions were within 5 cm of truth, 93.3% within 10 cm, and 80% of BMI categories were correctly classified. All three misclassifications occurred in subjects within 3 BMI units of a WHO category boundary — the expected region of highest inherent uncertainty

Table 4. Informal Multi-Subject Testing Summary (n = 15)

Criterion	Result	Percentage
Height within 5 cm of truth	11 / 15	73.3%
Height within 10 cm of truth	14 / 15	93.3%
Weight within 10 kg of truth	9 / 15	60.0%
Weight within 15 kg of truth	13 / 15	86.7%
BMI category correctly classified	12 / 15	80.0%
Misclassifications within 3 BMI units of boundary	3 / 3	100% of errors

The web application was tested across four primary interaction states: empty state before upload, image loaded with predict and clear buttons, processing spinner overlay, and final output with BMI results and Gemini health plan. Figures 1 through 4 capture these states.

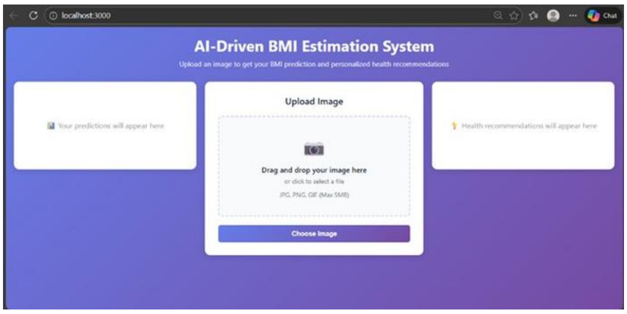

Figure 1 Web Application — Initial Empty State Before Upload

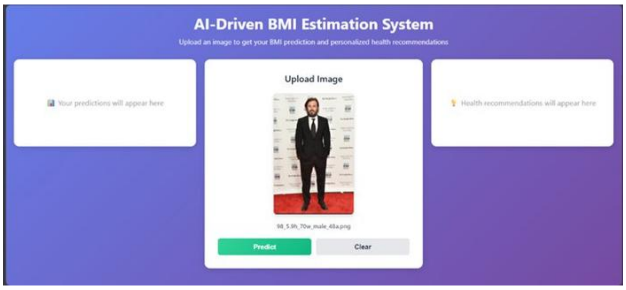

FIGURE 2 Web Application — Image Loaded with Predict and Clear Buttons

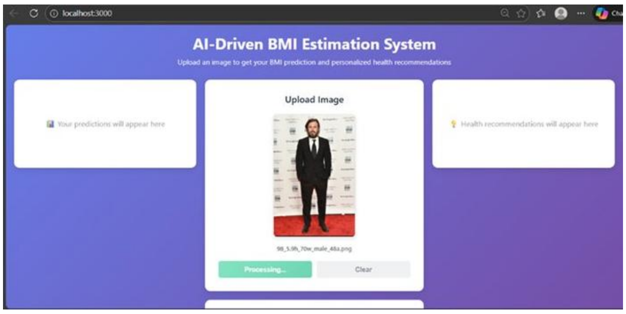

Figure 3 Web Application — Processing State with Spinner Overlay

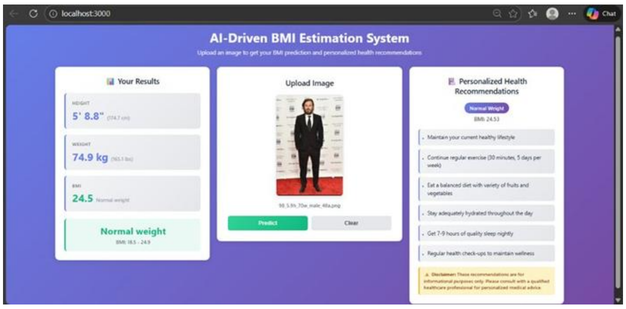

Figure 4 Web Application — Final Output: BMI with Health Plan

As seen in Figure 4, the system correctly predicted a height of 5'8.8" (174.7 cm) and weight of 74.9 kg, yielding a BMI of 24.5 classified as Normal weight. The left panel renders the biometric results with a colour-coded category badge and BMI scale bar; the right panel displays the Gemini-generated personalised health plan along with a mandatory health disclaimer. The mobile application was tested on two physical Android devices (Samsung Galaxy

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 04 April 2026
Page No: 2165-2171
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0289

A53, Android 13 and Redmi Note 11, Android 12) connected to the FastAPI backend over a shared mobile hotspot Figures 5 through 8 show the four mobile interface states.

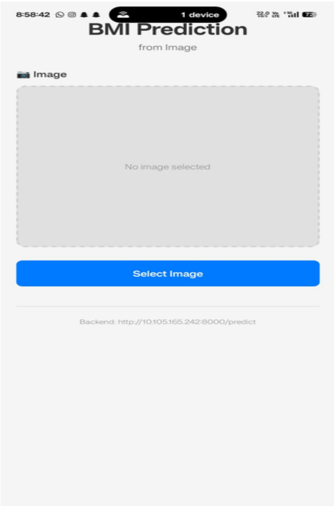

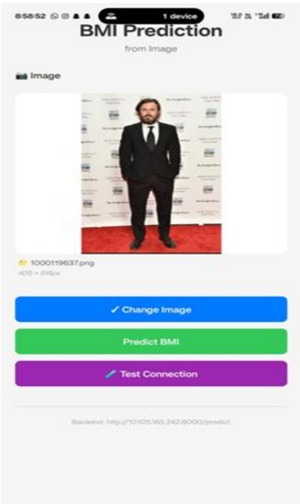

Figure 5 **Mobile — Empty State**
Figure 6 **Mobile — Image Selected**

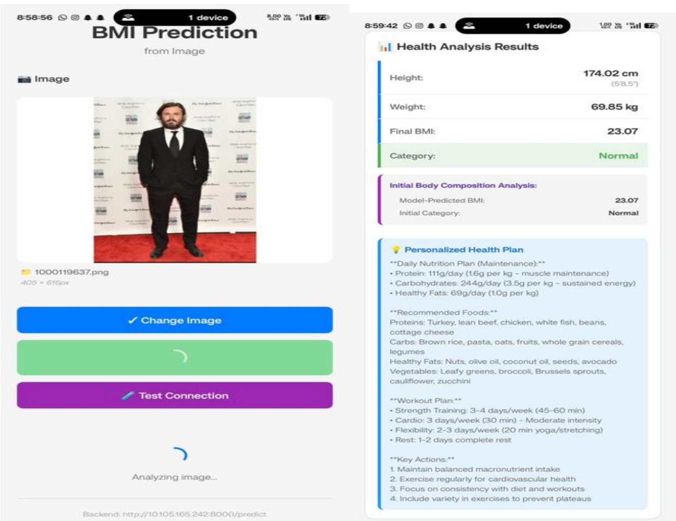

Figure 7, 8: Mobile

2.2.Discussion

The height MAE of 6.136 cm arises from three primary sources: perspective distortion (the model cannot access camera focal length or subject distance and must rely on population statistics for implicit scale assumptions), IMDB-Wiki demographic bias towards taller-than-average subjects causing regression-toward-the-mean at height extremes, and natural stance variation due to footwear and posture. The weight MAE of 9.8 kg reflects the fundamental ill-posedness of estimating a three-dimensional physical property from a 2D projection — body depth, muscle-fat composition, and bone density can vary by 10–20 kg between visually similar subjects.

Despite the weight MAE, the system remains practically useful for BMI category screening. For a person of 170 cm height, a weight error of 9.8 kg corresponds to a BMI error of approximately 3.4 units. Subjects more than 5 BMI units from the nearest category boundary — the majority of the population at the extremes — are unlikely to be misclassified. The Keypoint R-CNN quality gate rejected 11 of 47 submitted images (23.4%) during informal testing, with rejection reasons including feet not visible (6 cases), upper body partially cut off (3 cases), and low person detection confidence (2 cases). Forced inference on rejected images produced an average height error of 22.4 cm — confirming that the gate delivers a 3.6× reduction in prediction error for real-world images.

Gemini-generated recommendations were assessed qualitatively across ten invocations spanning all four BMI categories. Macronutrient targets were consistently personalised per kilogram of body weight, recommended foods reflected mainstream nutritional science, and exercise recommendations appropriately distinguished cardiovascular, resistance, and flexibility components by BMI category. All five end-to-end testers rated the Gemini recommendation as more useful than a generic search result, citing personalised macronutrient targets and exercise specificity as the most valued elements.

A comparison of the two deployment platforms is presented in Table 5. The web application benefits from Gemini-generated personalised recommendations and a visual BMI scale bar, while the mobile application offers lower end-to-end latency and gallery-native image picking with no drag-and-drop requirement.

Table 5 **Web vs. Mobile Application — Feature Comparison**

Aspect	Web Application	Mobile Application
Image selection	Drag-and-drop + file picker	Device gallery (expo-image-picker)
End-to-end latency	~2.8 s	~1.3 s
Recommendation	Gemini-generated (personalised)	Local fallback (category-level)
BMI visualisation	CSS gradient scale bar	Text-only display
Distribution	Browser URL (any platform)	Side-loaded APK (Android)

Conclusion

This paper has presented an AI-driven system that estimates BMI from a single full-body photograph and delivers personalised health recommendations via a large language model, deployed simultaneously as a web application and a native Android application. The system achieved a height MAE of 6.136 cm, a weight MAE of 9.8 kg, and 80% correct BMI category classification across 15 informal test subjects, with all misclassifications occurring at WHO category boundaries where measurement uncertainty is inherently highest. The Keypoint R-CNN quality gate reduced real-world height prediction error by 3.6× compared to unfiltered inference, demonstrating its practical value for user-facing deployments. The dual-platform architecture confirmed that a single FastAPI inference service can support both browser and mobile clients with identical numerical predictions in under three seconds.

The primary limitation is the weight MAE of 9.8 kg, inherent to the challenge of estimating a 3D property from a 2D image. Future work includes integrating Gemini directly into FastAPI so mobile users also receive personalised recommendations, containerising the backend for cloud deployment with HTTPS, retraining models on a more demographically diverse dataset to reduce IMDB-Wiki bias, and exploring monocular depth estimation as an additional input channel to improve weight prediction accuracy. Health equity — ensuring basic health information is accessible regardless of geography or income — is a defining challenge of global public health, and this system demonstrates a working proof of concept that places a credible, personalised health screening tool in the hands of anyone with a smartphone.

Acknowledgements

I would like to express my sincere gratitude to my faculty mentors and the Department of Information Technology at Kamaraj College of Engineering and Technology for their guidance, encouragement, and support throughout the development of the AI-Driven BMI Estimation And Health Recommendation System project. I also acknowledge the help of online communities and open-source contributors for providing valuable resources and tools. No external financial support was received for this work.

References

[1]. World Health Organization. (2023). Obesity and overweight. WHO Global Health Observatory. https://www.who.int/news-room/fact-sheets/detail/obesity-and-overweight

[2]. Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. Medical Image Computing and Computer-Assisted Intervention (MICCAI), 234–241. doi:10.1007/978-3-319-24574-4_28

[3]. He, K., Gkioxari, G., Dollár, P., & Girshick, R. (2017). Mask R-CNN. Proceedings of the IEEE International Conference on Computer Vision (ICCV), 2961–2969. doi:10.1109/ICCV.2017.322

[4]. Cao, Z., Simon, T., Wei, S.-E., & Sheikh, Y. (2017). Realtime multi-person 2D pose estimation using part affinity fields. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 7291–7299.

[5]. Caruana, R. (1997). Multitask learning. Machine

Learning, 28(1), 41–75. doi:10.1023/A:1007379606734
[6]. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521(7553), 436–444. doi:10.1038/nature14539
[7]. Velardo, C., & Dugelay, J.-L. (2010). Weight estimation from visual body appearance. Proceedings of the 4th IEEE International Conference on Biometrics: Theory, Applications and Systems (BTAS), 1–6. doi:10.1109/BTAS.2010.5634540
[8]. Jin, Z., Huang, J., Xiong, A., Pang, Y., Wang, W., & Ding, B. (2022). Attention guided deep features for accurate body mass index estimation. Pattern Recognition Letters, 154, 22–28. doi:10.1016/j.patrec.2021.12.015
[9]. Mana, N. A. M. A., Fook, C. Y., Chin, L. C., Vijean, V., Ardeenawatie, S., & Muthusamy, H. (2022). Deep learning-based body mass index (BMI) prediction using pre-trained CNN models. Advances in Intelligent Systems and Computing, 1370, 495–504. doi:10.1007/978-981-16-2123-9_48
[10]. Yousaf, N., Hussein, S., & Sultani, W. (2021). Estimation of BMI from facial images using semantic segmentation based region-aware pooling. Computers in Biology and Medicine,
[11]. 133, 104392. doi:10.1016/j.compbiomed.2021.104392
[12]. Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. Nature Medicine, 25(1), 44–56. doi:10.1038/s41591-018-0300-7
[13]. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 770–778. doi:10.1109/CVPR.2016.90
[14]. Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., & Zitnick, C. L. (2014). Microsoft COCO: Common objects in context. European Conference on Computer Vision (ECCV), Lecture Notes in Computer Science, 8693, 740–755. doi:10.1007/978-3-319-10602-1_48
[15]. Suri, B., Taneja, S., Bhanot, I., Sharma, H., & Raj, A. (2022). Cross-platform empirical analysis of mobile application development frameworks: Kotlin, React Native and Flutter. Proceedings of the International Conference on Information Management & Machine Intelligence (ICIMMI 2022). doi:10.1145/3590837.3590897
[16]. Adeshina, A. A. (2022). Building Python Web APIs with FastAPI: A fast-paced guide to building high-performance, robust web APIs. Packt Publishing. ISBN: 9781801076630