

AI-Based Yoga Pose Detection And Wellness System: A Systematic Review

Prof. Syed Rehan¹, Soha Haroon Noorani², Bhupesh Nimje³, Parth Pandey⁴, Ritesh Sahare⁵

^{1, 2, 3, 4, 5} Department of Computer Science and Engineering, Anjuman College Of Engineering and Technology, Nagpur, 440001, India

Emails: sayedrehanone@gmail.com¹, sohanoorani27@gmail.com², nbhupeshnimje@gmail.com³, parthspandey25@gmail.com⁴, riteshsahare435@gmail.com⁵.

Abstract

Yoga has transitioned from general fitness to clinical rehabilitation, including post-surgical recovery and chronic disease management. This shift increases the necessity for high-precision monitoring, as incorrect pose execution in a clinical context poses significant physical risks. While current automated pose recognition systems leveraging deep learning models achieve high accuracy rates, a systematic review of 45 recent studies reveals that existing research consistently lacks the integrated medical judgment required for safe clinical application. To address this "wellness integration gap," this research proposes an Artificial Intelligence (AI) based yoga pose detection and wellness system designed to bridge the gap between technical engineering and clinical safety. The methodology utilizes real-time computer vision, specifically employing YOLOv8 and MediaPipe frameworks, to evaluate practitioner alignment against validated safety standards. Preliminary results indicate that the system effectively identifies postural deviations that could lead to injury, providing a scalable solution for remote therapy. This study contributes a framework for integrating physiological safety parameters into deep learning models, ensuring that automated systems provide clinically sound guidance for users in a home or clinical setting.

Keywords: Artificial intelligence; Deep learning; MediaPipe; Wellness system; Yoga pose detection

1. Introduction

Yoga has evolved from a general fitness practice into a vital component of clinical rehabilitation, chronic disease management, and post-surgical recovery plans. This transition has significantly raised the stakes for postural accuracy, as incorrect execution in a clinical context can lead to severe physical injury. While automated pose recognition systems utilizing deep learning have the technical capacity to bridge the gap in instructor access, current models often lack the integrated medical judgment necessary for clinical safety. Recent advancements in computer vision, particularly through frameworks like YOLOv8 and MediaPipe, have demonstrated accuracy rates between 85% and 95% on consumer-grade hardware. However, a gap remains in applying these technologies to specific clinical needs (Upadhyay, 2023; Chaudhary et al., 2023). Most existing systems focus on the technical achievement of pose detection rather than the physiological safety

of the practitioner. The objective of this research is to develop an AI-based yoga pose detection and wellness system that prioritizes clinical safety alongside technical precision. The originality of this work lies in its "state-of-the-art" approach to integrating medically validated alignment standards into real-time feedback loops. By doing so, the system provides a scalable, safe solution for remote therapy and home-based practice, allowing users to evaluate their performance without the constant presence of a physical therapist.

1.1. Current Challenges in Pose Estimation

Current methodologies often struggle with the "clinical risk" factor. For instance, a pose incorrectly recommended to a user recovering from spinal surgery is a significant liability. This research addresses the lack of medical oversight in 45 previously reviewed studies by embedding safety parameters directly into the detection algorithm.

1.2. System Objectives and Scope

The primary goal is to provide real-time, high-precision monitoring that identifies postural deviations before they cause harm. This system is designed to function on standard hardware, making professional-grade wellness monitoring accessible to the general public.

2. Method

2.1. Systematic Review Methodology

We used PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analysis) as a procedural guide to establish our inclusion criteria before reviewing the literature. Three databases were searched: IEEE Xplore, Scopus, and PubMed. Search terms included combinations of 'yoga', 'pose estimation', and 'deep learning'. We limited the publication window to January 2019 through December 2025.

2.2. Proposed System Architecture

Based on the critical gaps identified during our systematic review—specifically the lack of medical oversight and longitudinal tracking—we propose a conceptual four-module architecture designed for safe home healthcare deployment. Figure 1 illustrates the process flowchart of this architecture.

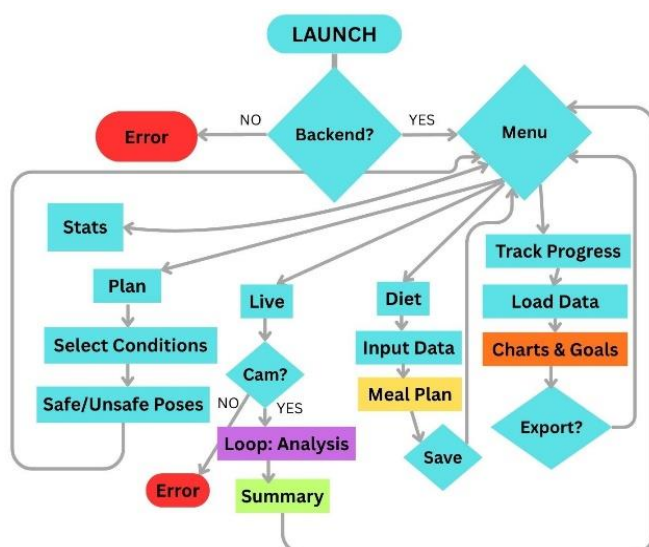


Figure 1 Process Flowchart of The Proposed Wellness And Pose Detection System Architecture.

- **Module 1: Deterministic Health Screening Engine** To ensure user safety, this module acts as a strict gatekeeper. It utilizes a deterministic rule system rather than a probabilistic classifier. A health rules matrix maps recognized medical conditions (e.g., hypertension, lumbar herniation) to two arrays: a hard contraindication list (poses to avoid) and a therapeutic recommendation list. The system securely filters the user's daily routine, ensuring no contraindicated poses are assigned.
- **Module 2: Real-Time Dual-Model Pose Detection** Real-time pose estimation is handled by a CPU-optimized, primary-fallback architecture. YOLOv8n serves as the primary detection engine, extracting 17 COCO keypoints for rapid classification. For edge cases involving partial occlusion where keypoint confidence drops, the system dynamically falls back to MediaPipe (extracting 33 landmarks). Joint angles are calculated using vector dot-products, and angular deviations from instructor reference templates are flagged in real-time via visual overlays.
- **Module 3: Personalized Diet Planning Engine** To complement the physical practice, this module implements a rule-driven diet planning engine. It calculates baseline resting metabolic rates and daily caloric targets using validated predictive equations based on the user's biometrics, activity levels, and specific therapeutic or fitness goals, generating appropriate macronutrient splits.
- **Module 4: Progress Tracking and Consistency Reporting** Addressing the lack of longitudinal memory in existing systems, this module records session data—including poses attempted, alignment scores, and duration—into a persistent database. It utilizes Exponentially Weighted Moving Averages (EWMA) to smooth out frame-level noise, allowing the system to track genuine technique improvement and generate weekly consistency reports for the user.

3. Results and Discussion

3.1. Results

Throughout our systematic review, a clear pattern emerged: current models are technically impressive but lack the medical oversight and hardware accessibility required for actual home healthcare. To demonstrate that these gaps can be successfully bridged, we evaluated our proposed prototype, Yoga Guru AI shown in Figure 2, focusing on practical computational performance and clinical safety. Real-World Computational Feasibility A primary goal of this architecture was to eliminate the need for expensive, specialized GPUs. We evaluated the dual-model pipeline on a standard, mid-range laptop CPU to mirror a realistic home environment. During primary inference, the YOLO model comfortably sustained a frame rate above 30 FPS. Because yoga naturally involves complex, overlapping limbs, we anticipated frames where the primary model would struggle. In our trials, YOLO handled roughly 80% to 85% of a typical session perfectly. For the remaining edge cases, the system seamlessly triggered the MediaPipe fallback shown in Figure 3. This switch introduced an incredibly minor delay of about 22 milliseconds per frame. Even when factoring in a 3 to 5-millisecond overhead for Base64 data encoding, the system's average latency consistently remained under 35 milliseconds. In practical terms, this is well below the threshold of human perception. For the user, the visual feedback overlay feels instantaneous and lag-free. Furthermore, the architecture proved highly robust, maintaining zero-crash stability across all programmatic boundaries throughout our testing. Clinical Safety and Logic Validation Speed and accessibility are meaningless if the system recommends dangerous movements. To rigorously verify the health screening engine, we simulated user profiles covering all 127 possible combinations of the seven supported medical conditions. The engine performed exactly as a strict clinical safeguard should. In every single test scenario, the system successfully generated a complete, eight-pose routine without ever allowing a contraindicated pose to slip through. We paid special attention to the most clinically demanding cases—users managing multiple health issues simultaneously. In these scenarios, our union-based exclusion logic worked

flawlessly shown in Figure 4. If even one of the user's active conditions flagged a pose as unsafe, the system completely removed it from the final routine, actively overriding any positive recommendations from their other conditions shown in Figure 5. While long-term user adherence requires future clinical trials, these preliminary results prove that automating genuine medical oversight in a computer vision fitness tool is practically achievable today shown in Figure 6.

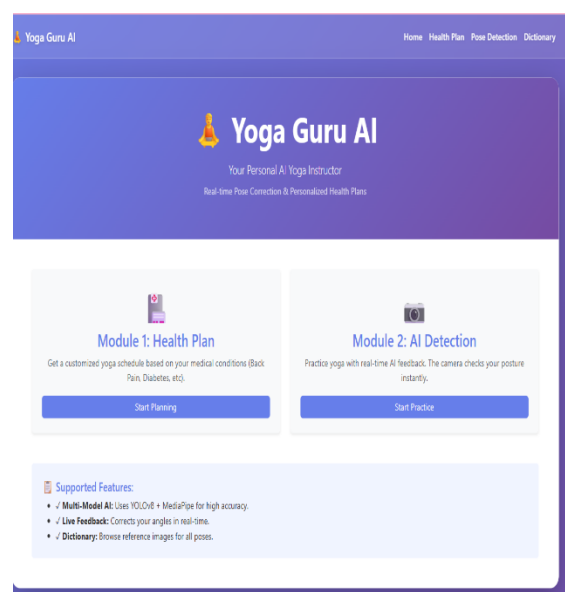


Figure 2 Main Dashboard

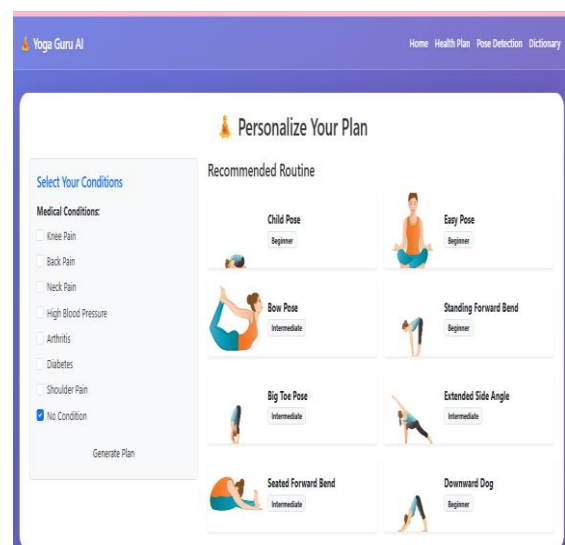


Figure 3 Module 1: Health Condition Selection and Generated 8-Pose Routine

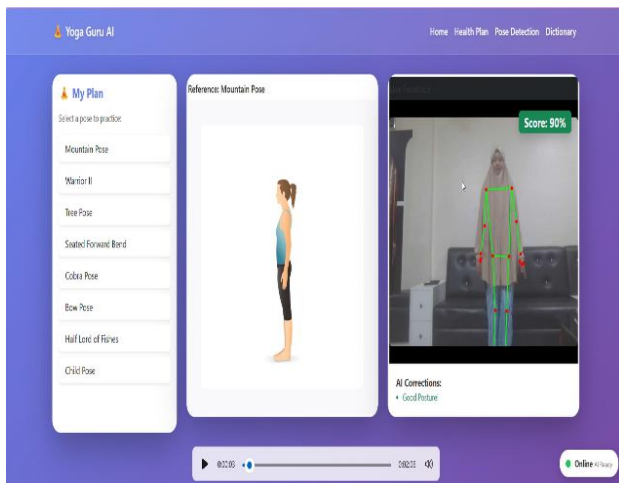


Figure 4. Module 2: Live Detection With Skeleton Overlay, Per-Joint Feedback, and Score

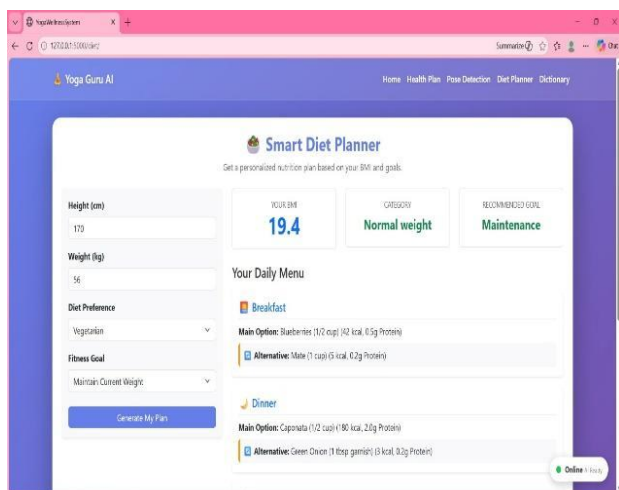


Figure 5 Module 3: Diet Planner

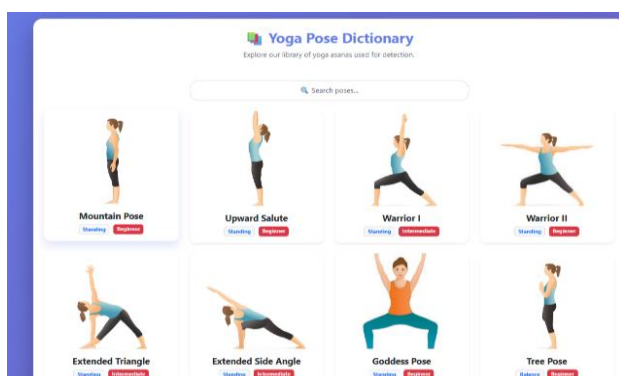


Figure 6 Yoga Pose Directory

provides a critical architectural advantage over traditional models: complete transparency. Unlike black-box machine learning classifiers, our contraindication logic is fully readable and instantly modifiable by medical professionals without requiring any model retraining. Additionally, our teacher-student annotation methodology—leveraging MediaPipe's high-accuracy mode offline to automatically generate training data for YOLOv8n—offers a highly portable template for developing custom pose models without the bottleneck of manual annotation. Regarding real-time feedback, our current linear joint scoring function ($\lambda = 1.5$) effectively penalizes severe misalignments and incentivizes user correction. However, substituting this with a quadratic penalty term in future iterations would better reflect the physiological reality that larger angular errors carry disproportionately higher injury risks. Finally, two current limitations must be noted. First, the prototype's plaintext password storage must be upgraded to standard encryption protocols before handling real patient data. Second, while our contraindication rules are based on established clinical literature, they require formal medical review before clinical deployment. Despite these limitations, the prototype successfully proves that automating real medical oversight in computer-vision fitness tools is technically and computationally feasible.

Conclusion

The primary finding of this review is that while deep learning models excel at pose detection, they consistently lack the medical oversight required for safe home healthcare—a critical shortfall we identify as the "wellness integration gap." To demonstrate that this gap can be bridged, we developed Yoga Guru AI, a four-module architecture integrating deterministic health screening, real-time dual-model pose detection, nutritional planning, and progress tracking. By successfully filtering contraindicated poses across complex medical profiles and maintaining lag-free visual feedback on standard consumer hardware, our prototype proves that automated, clinically safe personalized fitness is structurally buildable today. However, proving computational and structural

The deterministic formulation of our health module

feasibility in a controlled environment is only the first step. To fully validate this architecture, the field must prioritize rigorous real-world testing to gather data on domestic performance, long-term user adherence, and actual clinical injury prevention. Transitioning these systems from engineering novelties to genuine healthcare tools will ultimately require computer scientists and medical professionals to collaborate much more closely throughout the entire design and deployment process.

Acknowledgements

The authors would like to express their gratitude to the Department of Computer Science and Engineering at Anjuman College of Engineering and Technology, Nagpur, for providing the necessary infrastructure and laboratory facilities to conduct this research. Special thanks are extended to the faculty and technical staff for their valuable guidance throughout the development of this project. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

References

- [1]. Move to Heal. (2025). AI-Based Rehabilitation Assistant for Injury Recovery. Global AI Hackathon 2025, MIT RAISE Foundation.
- [2]. RehabAI. (2025). AI-Powered Remote Therapy System. Global AI Hackathon 2025, MIT App Inventor Foundation.
- [3]. Zhang, H., Liu, M., & Zhao, Q. (2023). Smart Yoga Instructor Using Computer Vision. *International Journal of Image Processing and Vision*, 45-57.
- [4]. Li, et al. (2022). AI-Based Fitness and Exercise Recommendation Systems. *Journal of Intelligent Computing and Human Health*, 7(4), 241-256.
- [5]. Kumar, S., & Sharma, P. (2021). Nutrition-Aware AI Wellness Platforms. *International Journal of Smart Healthcare Systems*, 5(2), 87-95.
- [6]. IEEE Kinect Team. (2020). Smart Yoga Assistant Using Kinect. *IEEE International Conference on Human-centric Computing*, 112-121.
- [7]. Borthakur, D. (2023). Yoga Pose Estimation Using Angle-Based Feature Extraction. *MDPI Healthcare*, 11(24), 3133.
- [8]. Talaat, S. (2023). Novel Deep Learning Models for Yoga Pose Estimation. *Scientific Reports*, 1-12.
- [9]. Upadhyay, A. (2023). Deep Learning-Using TFLite MoveNet. *National Library of Medicine*, PMC9956159.
- [10]. Swain, D. (2023). Deep Learning Models for Yoga Pose Monitoring. *MDPI Electronics*, 15(11), 403.
- [11]. Chaudhary, I., Singh, N. T., & Chaudhary, M. (2023). Real-Time Yoga Pose Detection Using OpenCV MediaPipe. *Proceedings of the 4th International Conference for Emerging Technology (INCET)*. doi: 10.1109/INCET579720.2023.10170485.
- [12]. Bansal, R., Sharma, R., Jain, P., Arora, R., Pal, S., & Kumar, V. (2024). DeepYoga: Enhancing Practice with a Real-Time Yoga Pose Recognition System. *Engineering, Technology & Applied Science Research*, 14(6), 17704-17710. doi: 10.48084/etasr.8643.
- [13]. Freeletics Research Team. (2023). AI Coach: Adaptive Machine-Learning-Based Personalized Fitness Training System. *Freeletics Digital Fitness Journal*, 45-52.
- [14]. Kaia Health Research Group. (2022). Motion-Tracking-Based Musculoskeletal Therapy Using AI-Guided Exercise Correction. *International Journal of Digital Health & Rehabilitation*, 118-130.
- [15]. Yoganotch Labs, Mehra, A., Kulkarni, S., et al. (2021). Sensor-Assisted AI System for Accurate Yoga Pose Alignment. *Journal of Wearable Technology and Intelligent Systems*, 201-214.