Towards Equitable Biomechanics: Rethinking Vision-Based Biomechanics for Resource-Limited Settings with Resourceful Design Strategies

Rune Chi Zhao*
runechi.zhao@anu.edu.au

Xiuyuan Yuan*
xiuyuan.yuan@anu.edu.au

Abstract

Vision-based biomechanics has enabled valuable tools for movement analysis, supporting applications such as rehabilitation monitoring, gait assessment, and neurological evaluation. However, these advances have largely been developed and validated in well-resourced environments, limiting their accessibility in resource-limited settings (RLS) where infrastructure, connectivity, and technical capacity remain constrained. In this paper, we examine why many current vision-based biomechanics pipelines fail to translate effectively to such contexts, identifying key barriers and important ethical considerations. We then propose a selection of design strategies aimed at improving applications for RLS, advocating for greater resourcefulness and practicality. We aim to inspire future directions and highlight how vision-based biomechanics can be redesigned to operate under constrained conditions. Finally, we propose that progress in the field should be evaluated not only by technical performance, but by the ability of systems to deliver meaningful and equitable impact in the environments where they are most needed.

1. Introduction

Since its emergence as a scientific discipline, biomechanics has contributed significantly to society by leading to the measurement and interpretation of human movement, enabling a wide range of clinical applications. In general, biomechanics studies how forces, motion, and body structure interact during movement, allowing researchers and clinicians to move from visual observations to quantitative assessment [32]. In recent years, advances in computer vision (CV) and machine learning (ML) have significantly expanded the field of biomechanics. From these developments, vision-based biomechanics systems can now infer information such as spatiotemporal gait parameters, joint kinematics, and, in certain pipelines, dynamic quantities such as joint moments, ground reaction forces, and musculoskeletal loading from video or other modalities when combined with biomechanical models [6, 30]. These capabilities are already supporting movement analysis across many clinical applications, including rehabilitation monitoring, gait assessment after stroke, Parkinsonian gait analysis, and musculoskeletal evaluation [17, 22, 25, 28, 31]. Through these applications, vision-based biomechanics have enabled more scalable and practical clinical tools, offering significant societal benefits.

However, these benefits have not been equitably distributed. Most research on digital health and AI systems originates from well-resourced, high-income healthcare environments, while resource-limited settings continue to face significant constraints in infrastructure, workforce capacity, and implementation support [4, 18]. Vision-based biomechanics reflects this broader pattern: even systems described as low-cost or markerless often assume access to specialized hardware, stable connectivity, and technical expertise.

Thus, there are numerous challenges preventing the implementation of valuable vision-based biometrics systems in resource-limited settings (RLS). In this paper, we use RLS settings broadly to refer not only to low- and middle-income countries (LMIC), but also to rural, remote, underserved, and otherwise marginalized communities where healthcare access is constrained by limited infrastructure, workforce shortages, geography, poverty, or exclusion, including certain communities within high-income countries [15, 38]. In these environments, health systems often operate with constrained diagnostic and treatment capacity, shortages of specialised workers, limited training pathways, and underdeveloped digital infrastructure, all of which can restrict the adoption of new technologies [7, 23, 35]. However, these are also the settings that possess the greatest burden of disease and the most urgent need for improvements in quality of healthcare. For example, stroke remains a major source of death and disability worldwide, with lower-SDI countries carrying a disproportionate burden and poorer quality of care [10]. Similarly, children and adolescents, WHO and UNICEF have likewise emphasised the large global burden associated with developmental disabilities, especially in such lower resource settings [33]. Many

*Equal contribution.

of these conditions directly involve movement, function, or rehabilitation, making them highly relevant to biomechanical monitoring and analysis. Thus, whilst RLS would arguably benefit more from such vision-based biomechanics systems, these technologies often remain less accessible in these settings, reflecting a gap that is not solely geographic, but rooted in broader patterns of under-representation and structural inequity.

This raises a clear technological equity concern. If vision-based biomechanics continues to be developed primarily around narrow datasets, idealised workflows, and well-resourced infrastructure assumptions, then its expansion may only improve care where resources are already abundant, leaving resource-limited settings further behind. This concern is an example of the risk of disparity amplification; where performance, access, and governance all improve in high-resource contexts, whilst populations with the greatest unmet need remain excluded [16, 36, 39].

Thus, we strongly believe that strong actions must be taken to prevent this trajectory. By rethinking how biomechanical systems are designed and deployed, we believe that it is possible to create more resource-aware biomechanics systems that sacrifice a small amount of precision while significantly expanding accessibility. Rather than waiting for technologies to gradually diffuse into lower-resource contexts, the field should actively explore more creative and resourceful design strategies that enable biomechanics to operate in more constrained environments. Accordingly, this short paper has three main aims. Firstly, we aim to examine and understand why many current vision-based biomechanics systems fail to translate well to RLS. Secondly, we propose practical design strategies that make such systems more realistic for deployment in those environments. Finally, we advocate for a shift in emphasis from benchmark performance alone towards implementation design, suitability, and equitable impact.

2. Current challenges of implementing vision-based biomechanics in RLS

There are several major challenges that prevent vision-based biomechanics from being deployed in RLS, compared to more resource-rich environments. Understanding these barriers is essential, as they can provide insights into what changes are required to reduce this gap. In this section, we examine a selection of these challenges to motivate more resourceful design strategies for future work.

2.1. Assumptions about infrastructure

Many current vision-based biomechanics pipelines still depend on infrastructure that is easier to provide in well-resourced laboratories than in routine care environments. OpenCap, for example, shows that smartphone video can be used to estimate human movement dynamics, but its validated workflow relies on two or more iOS devices and cloud-based processing to recover kinematics and dynamics [32]. Pose2Sim similarly bridges computer vision and biomechanics through multi-view video, triangulation, and OpenSim-based inverse kinematics, but assumes synchronized and calibrated camera setups with controlled viewpoints [26, 27]. In the broader literature, higher-precision systems often depend on multiple cameras, high frame rates, stable capture volumes, and substantial downstream compute, which are reasonable for laboratory validation, but inaccessible for many other settings [29, 34].

In RLS, connectivity alone can be a major bottleneck. ITU reports that in 2023, only 17% of people living in rural areas of low-income countries used the internet, and only 39% of the population in low-income countries had access to 4G coverage, with 5G nearly absent [14]. This is a major issue as many modern pipelines rely on stable uploads, cloud inference, or frequent software updates. In addition, hardware requirements such as stable camera placement, sufficient lighting, or consistent recording angles may also be difficult to guarantee.

2.2. Complex operational requirements

Many vision-based biomechanics systems also carry complex operational requirements that make real-world deployment difficult, even when the required hardware is available. For instance, reviews of current markerless motion-capture methods repeatedly report sensitivity to factors such as occlusion, clothing variation, and viewpoint shifts, as well as the propagation of upstream pose errors into downstream biomechanical estimates [5, 34]. Pagnon et al. further show that errors in 2D pose localization propagate into triangulation and then into inverse-kinematics outputs, which means that small failures early in the pipeline can distort the final biomechanical measures [26]. These issues are usually manageable in controlled environments with trained operators, but much less so in busy community settings. A workflow that requires careful calibration, controlled lighting, and technically informed quality control may be acceptable in a biomechanics lab, but it is much less practical for community health workers, non-specialist clinicians, physiotherapists, or families recording rehabilitation tasks at home. In their paper, Wade et al. explicitly note that current systems have been slow to transfer into biomechanics partly because they demand advanced computer-vision knowledge and programming skills [34]. Thus, successful deployment depends not only on model performance but also on training burden, usability, and ongoing support [7].

2.3. Dataset bias and limited representation

Finally, another challenge is that many of the datasets underlying vision-based biomechanics are unlikely to represent the communities in which these systems may eventually be deployed. Biomechanics validation studies and pose-estimation benchmarks are often built around adult participants, convenience samples, or structured movement settings [4, 19, 34]. Furthermore, in the biomechanics context, Wade et al. argue that current open-source pose-estimation datasets were never designed for biomechanical applications and are inconsistently labeled for biomechanically meaningful use.

In addition, models trained on predominantly Western, adult, or highly structured samples may not generalize well to populations with different body proportions, clothing styles, everyday movement patterns, or cultural practices. The broader computer vision literature shows why this is dangerous. Performance disparities across gender and skin tone have been well-documented, and attributed to how the training data does not adequately reflect the populations on which systems are later used, leading to much lower accuracy [2]. Recent fairness work in pose estimation has made similar conclusions, with recent studies noting that the field often lacks the demographic annotations needed even to measure whether performance differs systematically across groups [19, 37]. In healthcare, these gaps in data are not merely technical inconveniences, but can lead to systematic underperformance, misleading outputs, and unequal care, especially when deployed in already underserved settings [16]. It is important to note that whilst many forms of dataset bias, such as disparities across gender or skin tone, are not unique to resource-limited settings, their consequences may be amplified in these contexts. For many RLS, there is a severe lack of locally collected or annotated datasets, in additional to more limited opportunities for local validation, recalibration, or clinician oversight, meaning that such biases can more directly translate into harmful consequences like reduced system reliability and inequitable care. This concern aligns closely with the argument that AI systems for RLS must be evaluated not only for technical promise, but for whether they are ethically and operationally appropriate for the populations they are meant to serve [40]. Thus, for vision-based biomechanics, limited representation is not just a technical concern, but a real risk that can directly reinforce inequity if left unaddressed.

3. Resourceful design strategies for vision-based biomechanics in RLS

To respond to these challenges, we argue that more vision-based biomechanics systems should be designed explicitly with constrained environments in mind. Instead of focusing purely on performance, we propose that more vision-based biomechanics can consider design strategies that are most useful, realistic, and sustainable in such settings. Below, we outline several potential design directions that illustrate how more resource-efficient pipelines could expand access to biomechanical analysis in RLS.

3.1. Pose-First Biomechanics Pipelines

One promising strategy is to separate visual data collection from downstream biomechanical analysis by relying more heavily on pose-based representations than on raw video. Modern pose estimation systems such as OpenPose and BlazePose can extract skeletal keypoints from image frames and represent the body as joint coordinates rather than dense RGB imagery [1, 3]. From a practical perspective, this has several key benefits, it can reduce storage and transmission requirements, make on-device processing more realistic, and avoid long-term storage of identifiable video data. For RLS, those improvements could make a very meaningful difference. Applications that adopt a post-first pipeline may be far more compatible with intermittent connectivity and limited storage, compared to one that requires continuous upload of full-resolution video. Furthermore, pose-first workflows are also already compatible with existing biomechanics pipelines. For example, OpenCap already uses pose estimation as an intermediate step before estimating kinematics and dynamics, while newer work on marker enhancement shows that improving the mapping from video keypoints to biomechanically meaningful markers can further strengthen the downstream utility of pose-driven systems [9, 32].

Furthermore, pose-first designs may also be ethically preferable in some contexts, although this still needs to be carefully evaluated. Since skeleton representations remove many appearance-based visual cues, they may reduce some forms of bias and privacy risk relative to raw image pipelines. They may also reduce the risk of facial recognition, identity tracking, and secondary visual surveillance. However, pose data is not neutral as body shape, disability, age, and other sensitive attributes can still sometimes be inferred from motion and skeletal structure, meaning it is not an entirely risk-free representation [13, 24]. In healthcare settings, especially in resource-limited communities that may already be more vulnerable to extractive data practices, data governance and autonomy is extremely important and something that must be carefully considered. Thus, even with a pose-first system, it is still vital to ensure explicit governance, privacy regulations, and clear benefit for the communities providing the data.

3.2. Smartphone-First Biomechanics

Another promising direction is to utilise smartphones as a primary sensing platform. Modern smartphones are widely distributed and combine cameras, inertial measurement units, and increasingly capable processors in one device. Most importantly, mobile phones are often the main route to digital access in low- and middle-income countries. GSMA reports that, in 2023, more than 3.7 billion people in LMICs accessed the internet on a mobile phone, and that mobile accounted for 84% of broadband connections in these settings [11, 12]. Affordability of new systems remains a major barrier for RLS, and so designing applications around existing phones over building new dedicated capture systems can be a strategy for minimizing cost.

There are already strong case studies showing what smartphone-first biomechanics can look like. OpenCap is the clearest example, with Uhlrich et al. showing that synchronised smartphone videos can be used to recover three-dimensional kinematics and musculoskeletal dynamics with useful accuracy, while dramatically reducing cost and time compared with laboratory-based workflows. They also report a 100-subject field study in which a clinician used OpenCap to estimate musculoskeletal dynamics much faster and at far lower cost than a traditional lab-based approach [32]. This application suggests that smartphone-based systems are a credible potential option for vision-based biomechanics applications in RLS.

3.3. Low-Frame-Rate and Imperfect-Video Biomechanics

Another useful direction is to design systems that remain useful even when video quality is limited. Many pose-estimation and biomechanical pipelines are developed on relatively clean, high-frame-rate recordings, but low-cost devices or real-world telehealth recordings may not meet those conditions. However, this lower quality video data may still be valuable. Recently, Dunn et al. showed that frame interpolation can improve the precision of stance time, swing time, and step timing in markerless single-camera gait analysis, partially compensating for reduced temporal resolution, demonstrating that there are alterative strategies that can be tried [8].

However, there still may be an inevitable decrease in precision if lower frame-rates or imperfect videos are used. This raises an important design question for RLS applications: should a system be rejected because it is less precise, or should it still be used if it can provide clinically useful information where the alternative is no quantitative movement assessment at all? For example, an imperfect but robust low-frame-rate system may still support triage, screening, follow-up, or rehabilitation monitoring, especially if it is transparent about uncertainty and limited to measurements that remain reliable under sparse temporal sampling. Thus, it is important to note that the priority in different contexts may vary and lead to vastly different approaches.

3.4. Low-bandwidth and local-first processing architectures

As mentioned previously, connectivity constraints remain a major challenge in many RLS, where continuous uploading of high-resolution video may be impractical or too expensive [11, 14]. One approach is to adopt local-first processing architectures, in which pose estimation and biomechanical feature extraction occur directly on-device, and only compact skeletal trajectories or derived features are transmitted when connectivity is available. Recent models such as BlazePose demonstrate that real-time pose inference can run on mobile hardware, enabling edge-based motion analysis without cloud dependence [1]. Similar ideas have appeared in earlier work on skeleton-based compression and privacy-preserving sensing, which emphasize transmitting structured motion representations rather than raw video streams [20, 21].

3.5. Community-Based Deployment

Finally, deployable vision-based biomechanics in RLS should be designed around the people who will actually use it. In most cases, they will not be experts in biomechanics, but people like community health workers, physiotherapists, rehabilitation assistants, coaches, teachers, or family caregivers. Broader digital health literature shows that successful deployment depends heavily on usability, training burden, workflow fit, and community acceptance, rather than model performance alone [7, 23]. For biomechanics, this means systems should be simplified for non-specialist operation, and paired with local training. Furthermore, they should be designed with cultural and community considerations in mind, which would help build trust and connection with the community. Ultimately for a technology to be successful, it cannot just be deployed, instead it must be trusted, accepted, and then used.

4. Conclusion

By combining computer vision with biomechanical modeling, vision-based biomechanics has enabled transformative new tools for rehabilitation, gait analysis, and health monitoring. However, these benefits of these technologies remain unevenly distributed. Despite having the greatest need for such tools, RLS face great difficulty in accessing them, as they are often created with the infrastructure and expertise of more well-resourced environments in mind. Thus, in this paper, we argue that addressing this gap requires a shift toward more resource-aware approaches.

Ultimately, progress in vision-based biomechanics should be judged not only by benchmark performance, but by whether systems can be implemented responsibly and sustainably in the contexts they are intended to serve. In many cases, a slightly less precise system that is affordable, usable, and locally deployable may deliver far greater real-world impact than a technically superior system that remains inaccessible. Thus, if vision-based biomechanics is to realize its full societal potential, it must be designed not only to perform well in ideal conditions, but to function effectively where it matters most.

References

[1] Valentin Bazarevsky, Ivan Grishchenko, Karthik Raveendran, Tyler Zhu, Fan Zhang, and Matthias Grundmann. Blazepose: On-device real-time body pose tracking. *arXiv preprint arXiv:2006.10204*, 2020. 3, 4

[2] Joy Buolamwini and Timnit Gebru. Gender Shades: Intersectional accuracy disparities in commercial gender classification. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency (FAT*)*, pages 77–91. PMLR, 2018. 3

[3] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2D pose estimation using part affinity fields. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1302–1310, 2017. 3

[4] Leo Anthony Celi, Jacqueline Cellini, Marie-Laure Charpignon, and et al. Sources of bias in artificial intelligence that perpetuate healthcare disparities—A global review. *PLOS Digital Health*, 1(3):e0000022, 2022. 1, 3

[5] R. James Cotton, Allison DeLillo, Anthony Cimorelli, Kunal Shah, J. D. Peiffer, Shawana Anarwala, Kayan Abdou, and Tasos Karakostas. Markerless motion capture and biomechanical analysis pipeline. *arXiv preprint arXiv:2303.10654*, 2023. 2

[6] Scott L. Delp, Frank C. Anderson, Allison S. Arnold, and et al. Opensim: Open-source software to create and analyze dynamic simulations of movement. *IEEE Transactions on Biomedical Engineering*, 54(11):1940–1950, 2007. 1

[7] Israel Júnior Borges do Nascimento, Hebatullah Abdulazeem, Lenny Thinagaran Vasanthan, and et al. Barriers and facilitators to utilizing digital health technologies by healthcare professionals: An overview of systematic reviews. *npj Digital Medicine*, 6:161, 2023. 1, 2, 4

[8] Marcus Dunn, Adam Kennerley, Zhane Murrell-Smith, Kate Webster, Kane Middleton, and Jon Wheat. Application of video frame interpolation to markerless, single-camera gait analysis. *Sports Engineering*, 26:22, 2023. 4

[9] Antoine Falisse, Scott D. Uhlrich, Akshay S. Chaudhari, Jennifer L. Hicks, and Scott L. Delp. Marker data enhancement for markerless motion capture. *IEEE Transactions on Biomedical Engineering*, 72(6):2013–2022, 2025. 3

[10] GBD 2019 Stroke Collaborators. Global, regional, and national burden of stroke and its risk factors, 1990–2019: A systematic analysis for the Global Burden of Disease Study 2019. *The Lancet Neurology*, 20(10):795–820, 2021. 1

[11] GSMA. The state of mobile internet connectivity 2023, 2023. 4

[12] GSMA. The state of mobile internet connectivity 2024, 2024. 4

[13] Carlos Hinojosa, Juan Carlos Niebles, and Henry Arguello. Learning privacy-preserving optics for human pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2573–2582, 2021. 3

[14] International Telecommunication Union. Facts and figures 2023, 2023. 2, 4

[15] K. Johnston and et al. Barriers and facilitators to the implementation of digital health services in rural and remote settings: A systematic review. *Journal of Medical Internet Research*, 2020. 1

[16] N. Joseph and et al. Algorithmic bias in public health AI: A threat to equity. *Public Health*, 2025. 2, 3

[17] J. Kim, R. Kim, K. Byun, N. Kang, and K. Park. Assessment of temporospatial and kinematic gait parameters using human pose estimation in patients with parkinson's disease: A comparison between near-frontal and lateral views. *PLOS ONE*, 20(1):e0317933, 2025. 1

[18] Alain B. Labrique. Ensuring digital health interventions are fit for purpose in low-resource settings. *npj Digital Medicine*, 3:146, 2020. 1

[19] Julienne LaChance, William Thong, Shruti Nagpal, and Alice Xiang. A case study in fairness evaluation: Current limitations and challenges for human pose estimation. In *AAAI Workshop on Representation Learning for Responsible Human-Centric AI (R2HCAI)*, 2023. 3

[20] X. Liang and et al. Skeleton-based privacy-preserving smart activity sensor for senior care and patient monitoring. *Preprints*, 2024. 4

[21] J. M. Lien and et al. Skeleton-based data compression for multi-camera tele-immersion system. In *Advances in Multimedia Information Processing – PCM 2007*. Springer, 2007. 4

[22] Luca Lonini, Yaejin Moon, Kyle Embry, R. James Cotton, Kelly McKenzie, Sophia Jenz, and Arun Jayaraman. Video-based pose estimation for gait analysis in stroke survivors during clinical assessments: A proof-of-concept study. *Digital Biomarkers*, 6(1):9–18, 2022. 1

[23] Shravya Mantena, Ananth Prasad, Saif Khairat, and et al. Improving community health-care screenings with smartphone AI technologies. *The Lancet Digital Health*, 3(2):e96–e97, 2021. 1, 4

[24] Md Mohibullah, Y. Suda, Y. Hironaka, T. Miyawaki, R. Suzuki, M. Hasan, and Y. Kobayashi. Privacy-preserving 3D human skeleton reconstruction from ankle-level 2D LiDAR using deep learning. *Neurocomputing*, page 131862, 2025. 3

[25] S. Natraj, T. Messmer, Y. Fujii, and et al. 3d pose estimation for scalable remote gait kinematics assessment. *npj Digital Medicine*, 9:37, 2026. 1

[26] David Pagnon, Mathieu Domalain, and Lionel Reveret. Pose2sim: An end-to-end workflow for 3D markerless sports kinematics—Part 1: Robustness. *Sensors*, 21(19):6530, 2021. 2

[27] David Pagnon, Mathieu Domalain, and Lionel Reveret. Pose2sim: An end-to-end workflow for 3D markerless sports kinematics—Part 2: Accuracy. *Sensors*, 22(7):2712, 2022. 2

[28] Jeffrey A. Reinbolt, Ajay Seth, Scott L. Delp, and et al. Simulation of human movement: Applications using opensim. *Procedia IUTAM*, 2:186–198, 2011. 1

[29] Sofia Scataglini, Eveline Abts, Cas Van Bocxlaer, Maxime Van den Bussche, Sara Meletani, and Steven Truijen. Accuracy, validity, and reliability of markerless camera-based 3D motion capture systems versus marker-based 3D motion capture systems in gait analysis: A systematic review and meta-analysis. *Sensors*, 24(11):3686, 2024. 2

[30] Ajay Seth, Jennifer L. Hicks, Thomas K. Uchida, and et al. Opensim: Simulating musculoskeletal dynamics and neuromuscular control to study human and animal movement. *PLOS Computational Biology*, 14(7):e1006223, 2018. 1

[31] Jesper Stenum, Michael M. Hsu, Albert Y. Pantelyat, and Ryan T. Roemmich. Clinical gait analysis using video-based pose estimation: Multiple perspectives, clinical populations, and measuring change. *PLOS Digital Health*, 3(3): e0000467, 2024. 1

[32] Sean D. Uhlrich, Antoine Falisse, Łukasz Kidziński, and et al. Opencap: Human movement dynamics from smartphone videos. *PLOS Computational Biology*, 19(10):e1011462, 2023. 1, 2, 3, 4

[33] UNICEF and World Health Organization. Global report on children with developmental disabilities, 2023. 1

[34] Logan Wade, Laurie Needham, Polly McGuigan, and James Bilzon. Applications and limitations of current markerless motion capture methods for clinical gait biomechanics. *PeerJ*, 10:e12995, 2022. 2, 3

[35] World Health Organization. *Recommendations on digital interventions for health system strengthening*. World Health Organization, Geneva, 2019. 1

[36] World Health Organization. *Ethics and governance of artificial intelligence for health*. World Health Organization, Geneva, 2021. 2

[37] Alice Xiang and et al. Fair human-centric image dataset for ethical AI benchmarking. *Nature*, 648(8092):97–108, 2025. 3

[38] Harini M. Yapa and Till Bärnighausen. Low resource settings as a contextual determinant of health. *Tropical Medicine & International Health*, 2018. 1

[39] H. Yu and et al. AI, global health, and ethical considerations: Risks of exacerbating inequality without localized governance and engagement. *Public Health*, 2024. 2

[40] Rune Chi Zhao and Xiuyuan Yuan. AI in healthcare for resource limited settings: An exploration and ethical evaluation. In *Companion Proceedings of the ACM Web Conference 2025*, 2025. 3