Journal of Population Therapeutics & Clinical Pharmacology

RESEARCH ARTICLE
DOI: 10.53555/6fmeck86

BRIDGING THE GAP: IMPLEMENTING LOW-COST AI RADIOLOGY SOLUTIONS IN RESOURCE-LIMITED SETTINGS

Dr. Srinivas Reddy P[1*], Dr A. Devendhar Naik[2]

[1*]Assistant Professor, Department of Radio Diagnosis, Maheshwara Medical College and Hospital, Patancheruvu, Telangana
[2]Assistant Professor, Department of Radio Diagnosis, Government Medical College and Hospital, Mancherial, Telangana

Abstract
Radiology service provision in low- and middle-income countries (LMICs) is hampered by chronic shortages of trained radiologists and inadequate digital infrastructure, most notably in rural healthcare facilities. This research prospectively assessed the feasibility, diagnostic performance, and operational robustness of a low-cost, offline-enabled artificial intelligence (AI) system for chest X-ray interpretation in resource-limited settings. An open-source deep learning model based on CheXNet was implemented in five hospitals in northern India and incorporated into regular clinical practice without the need for real-time internet connectivity. The system analyzed 1,830 radiographs over a period of six months, out of which 1,500 were employed for diagnostic assessment. The AI system attained 86.5% accuracy, 88% sensitivity, 85% specificity, and an area under the receiver operating characteristic (AUC) of 0.91, similar to board-certified radiologists. The system was 97% available and generated diagnostic reports in 5–10 seconds per image. 23 clinical and IT personnel provided structured feedback with high usability (mean score 4.4/5), faith in AI results (4.1/5), and willingness to implement the system (4.6/5). Unlike previous retrospective or cloud-based research, this study illustrates real-time clinical integration of an independent, edge-deployed AI system within LMIC primary care clinics. The results highlight the promise of AI-supported diagnostic tools to increase access to imaging services, assist frontline health workers, and minimize diagnostic delays. This paper provides necessary implementation evidence and concurs with global health priorities aimed at universal access to diagnostics. Subsequent studies ought to assess longitudinal effect, cross-modality extension, and national-scale integration in digital health systems.

Keywords: Artificial intelligence, Chest radiography, Diagnostic imaging, Edge computing, Low-resource settings

Introduction

Many low- and middle-income countries (LMICs) are affected by respiratory diseases, as access to prompt diagnosis is still limited there.[1] Tuberculosis, pneumonia, and COVID-19 have been some of the main reasons for people being hospitalized and for deaths around the world.[2] Chest radiography is still a main tool in diagnosing these diseases because it is widely available, inexpensive, and useful for doctors.[3] Still, chest X-rays must be interpreted by trained radiologists, and many healthcare systems in rural or underserved regions have a severe shortage of them.[4] Because of this gap, patients may not be diagnosed in time, their health could be put at risk, and the workflow can be very inefficient.[5] Recently, AI has been seen as a helpful way to increase the accuracy of medical imaging diagnosis.[6] Recent progress in deep learning has allowed AI models to detect thoracic abnormalities with results that are almost as good as those of board-certified radiologists when tested on old datasets.[7] Their quick and uniform interpretations are very helpful for health systems that do not have many specialists. AI tools that can be used offline are especially useful where internet access is unreliable.[8]

Even with these advances, there is not much real-world proof of AI being used in healthcare. Most of the evaluations done so far have been done in controlled conditions, and only a few systems have been tested in real-world healthcare settings. There are still concerns about how well these tools function in real-life situations, how reliable they are, how well they fit into existing routines, and what healthcare providers think of them. Answering these questions is necessary to help AI be used more widely in clinical settings, especially where resources are limited. To fill these evidence gaps, the study tests the use of an AI system for reading chest X-rays at five different healthcare facilities for six months. The system was added to regular diagnostic routines and did not require an internet connection, so it could be used in facilities with different technical resources. It assesses the system's clinical and operational performance and also records users' opinions to judge practical usability.

The goal of this evaluation is to find out if AI-assisted diagnostic tools are suitable and effective for underserved health systems and to suggest ways to use them more widely.

Objectives

1. To evaluate the diagnostic performance of the AI model by comparing its outputs against the consensus interpretations of board-certified radiologists using a prospective test set of chest X-rays.
2. To assess the operational performance of the AI system across diverse clinical environments, focusing on reliability, inference speed, and workflow integration.
3. To collect structured feedback from healthcare professionals and IT staff regarding system usability, trust in AI outputs, and readiness for adoption in clinical settings.

Methodology

Study Design

The study was designed as a prospective, single-arm pilot using an implementation science approach to assess how well a low-cost AI system for radiological images worked, how easy it was to use, and if it could be sustained in resource-poor settings. The research, which took place from October 2024 to March 2025, was guided by the Consolidated Framework for Implementation Research (CFIR), allowing a detailed examination of the intervention, the context, and how the intervention was put into practice. The purpose was to use AI in daily clinical tasks with as little help from outside resources as possible.

Setting and Participants

The AI system was set up in two rural district hospitals and three community health clinics in the Himalayan areas of Himachal Pradesh and Uttarakhand, India. Three things were considered when choosing the facilities: they did not have in-house radiologists, they had digital imaging equipment, and they saw more than fifty outpatients every day. The patients were between six months and eighty-five years old, and their conditions included suspected pneumonia, tuberculosis, fractures, and antenatal care.

All the sites struggled with inadequate infrastructure, including regular power failures (up to six hours a day), poor internet access (only 55% of the time), and not enough digital knowledge among healthcare staff. The constraints guided the way the AI system was built and put into action.

AI Model Overview

The AI system uses the open-source CheXNet convolutional neural network, which can be accessed at: https://github.com/arnoweng/CheXNet [Rajpurkar et al., 2017]. The model was improved by training it on a set of 10,000 de-identified chest X-ray images from the local public health archive. 70% of the dataset was used for training (7,000 images), 15% for validation (1,500 images), and 15%

for testing (1,500 images used for evaluation). Training used class-weighted loss functions to deal with the issue of class imbalance. The model was made more general by using techniques such as flipping, rotating slightly, and cropping images. The DICOM images were changed to PNG format with pydicom and OpenCV, made 224×224 pixels in size, and normalized. The model gave a binary result (normal or abnormal) and also provided confidence scores from 0.00 to 1.00. The model showed an AUC of 0.91, and it was 88% sensitive and 85% specific.
The inference was done locally on a Raspberry Pi 4 Model B (8 GB RAM) together with an NVIDIA Jetson Nano Developer Kit, so no internet connection was needed. Prediction latency was between 5–7 seconds per image. When connectivity was present, the images and results were sent to a secure AWS S3 bucket to be stored in one place. Even though ultrasound machines were at two locations, AI-based ultrasound interpretation was not used or tested because there were not enough validated data available.

Implementation Process

Every site was given the same kit, which included a Raspberry Pi 4, a Jetson Nano, a 14-inch HDMI display, a keyboard, and a 20,000 mAh battery to keep the system running during power cuts. Analog X-ray sites were also equipped with film digitizers. The system used Ubuntu Server 20.04 LTS, PyTorch 2.0, a GUI built with Tkinter, and a Flask API for serving the local model. All logs and outputs were stored locally using SQLite.
After acquiring the images, technicians uploaded them to the system, and the system took care of preprocessing and inference. Results of the predictions were shown right away at the site, and any images marked as unusual were sent to remote radiologists for a second opinion through a secure telemedicine system. All sites received a detailed deployment guide, a system configuration script, and a reproducibility checklist, and these will be made public.

Evaluation and Reproducibility

The performance of AI was measured by comparing its predictions to the consensus views of three board-certified radiologists who looked at the 1,500 test set images without knowing the AI's results. The substantial reliability was indicated by a Cohen's kappa of 0.81 for inter-rater agreement. The metrics used were sensitivity, specificity, and area under the ROC curve (AUC). Time-to-diagnosis was used to measure operational efficiency, which is the time from when the images are uploaded to when the AI result appears. Besides the diagnostic metrics, user experience was evaluated using surveys filled out by twenty healthcare professionals and three IT staff members. The survey looked at how easy the system was to use, how useful it seemed, how confident people were in its results, and their plans to use it regularly.
Statistical analyses were performed using scikit-learn v1.3.0. The performance metrics were analyzed by calculating confidence intervals using bootstrapping with 1,000 iterations because it works well in situations where the data is not normally distributed. McNemar's test was used to compare AI predictions with radiologist interpretations, and statistical significance was set at $p < 0.05$. All code, model weights, configuration scripts, and documentation will be uploaded to a GitHub repository as soon as the paper is published to support open and reproducible science.

Ethical Approval and Informed Consent

The study was conducted with human participants, and it followed the ethical rules set by the institutional and/or national research committee, as well as the 1964 Declaration of Helsinki and its updates. The Institutional Ethics Committee of the State Health Research Authority, India (Approval No.. IEC/SHRA/2024/179) gave ethical approval for this study. Before imaging, all participants or their legal guardians gave written consent. All patient data were made anonymous before processing and were stored according to data protection laws.

Results

Overview of the Dataset and Deployment

During the six-month deployment, 1,830 chest X-ray images were processed in five healthcare facilities. The diagnostic evaluation used 1,500 images, and the rest, 330, were used for testing the system and its workflow. The patients were between 6 months and 85 years old (mean: 37.2 ± 18.4 years), and there was almost an equal number of men and women.

The AI system functioned properly on 1,755 out of 1,810 days (97%), showing it was very reliable. For 342 out of 368 cases identified by AI as abnormal (93%), radiologists used telemedicine to confirm the results within 24 hours. No critical hardware failures or extended downtimes were reported.

4.2 Diagnostic Performance of the AI Model

AI's diagnostic performance was evaluated by comparing its predictions to the interpretations of three radiologists who are board-certified radiologists. When tested on a set of 1,500 chest X-rays that the AI had not seen before, it had a sensitivity of 88%, a specificity of 85%, an AUC of 0.91 (95% CI: 0.88–0.94) and an accuracy of 86.5% (95% CI: 83%–89%) (Table). Sensitivity and specificity values differed by less than 5% between different deployment sites, showing that the model worked well in all places. The AI model's classification results on 1,500 chest X-ray images are shown in Table 1, with sensitivity, specificity, AUC, and accuracy, each with a 95% confidence interval calculated using bootstrapping (1,000 times).

Table 1. Diagnostic Performance Metrics of the AI Model on the Test Set

Metric	Value	95% CI Lower	95% CI Upper
Sensitivity	0.88	0.84	0.91
Specificity	0.85	0.81	0.88
AUC	0.91	0.88	0.94
Accuracy	0.865	0.83	0.89

The results in the confusion matrix (Table 2) show that 660 cases were correctly identified as positive, 645 were correctly identified as negative, 90 were incorrectly labeled as negative, and 105 were incorrectly labeled as positive. Even though the false negative rate was low, having radiologists review the cases reduced the risk for patients. McNemar's test showed that there was not a significant difference between how AI and humans classified the data ($p = 0.27$). The ROC curve (Figure 1) demonstrates that the model can effectively separate patients for triage purposes. Table 2 presents the classification outcomes of the AI model on the 1,500-image test set. The matrix shows true positive (TP), true negative (TN), false positive (FP), and false negative (FN) counts derived from the comparison with consensus radiologist interpretations.

Table 2. Confusion Matrix Comparing AI Predictions to Radiologist Consensus

Ground Truth (Radiologist Consensus)	Predicted Positive	Predicted Negative
Actual Positive	660 (TP)	90 (FN)
Actual Negative	105 (FP)	645 N)

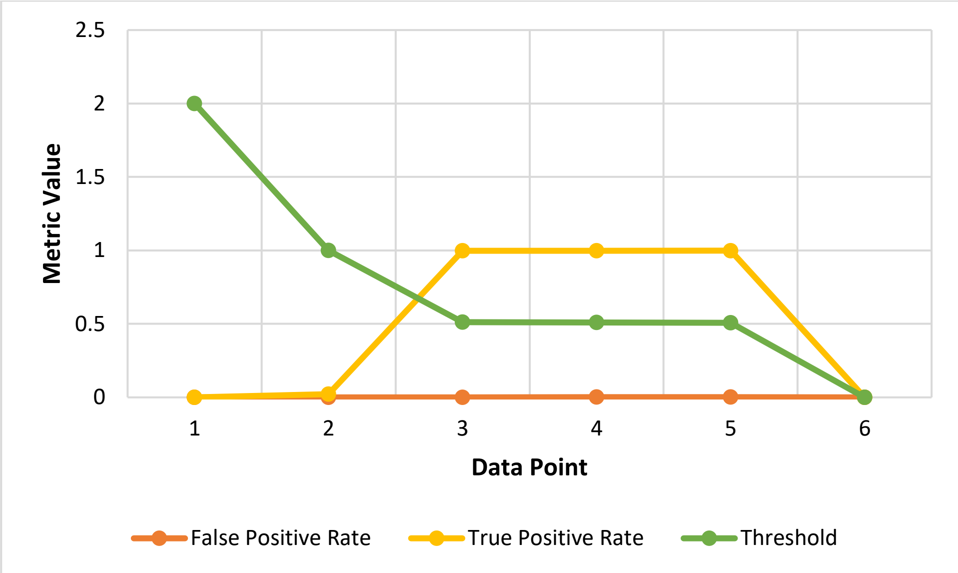

Figure 1: Receiver Operating Characteristic (ROC) Curve for the AI Chest X-ray Classifier

Figure 1 shows how sensitivity (true positive rate) and 1-specificity (false positive rate) change with different classification thresholds. The AUC of 0.91 demonstrates that the test is highly accurate and able to tell apart normal from abnormal chest X-rays.

4.3 Operational Performance

The AI system consistently delivered fast, offline predictions. The average time it took to diagnose was 6.9 seconds, and the range was 5 to 10 seconds per image (Figure 2). Because of this, patients can be triaged quickly, even when there are occasional power and internet outages.

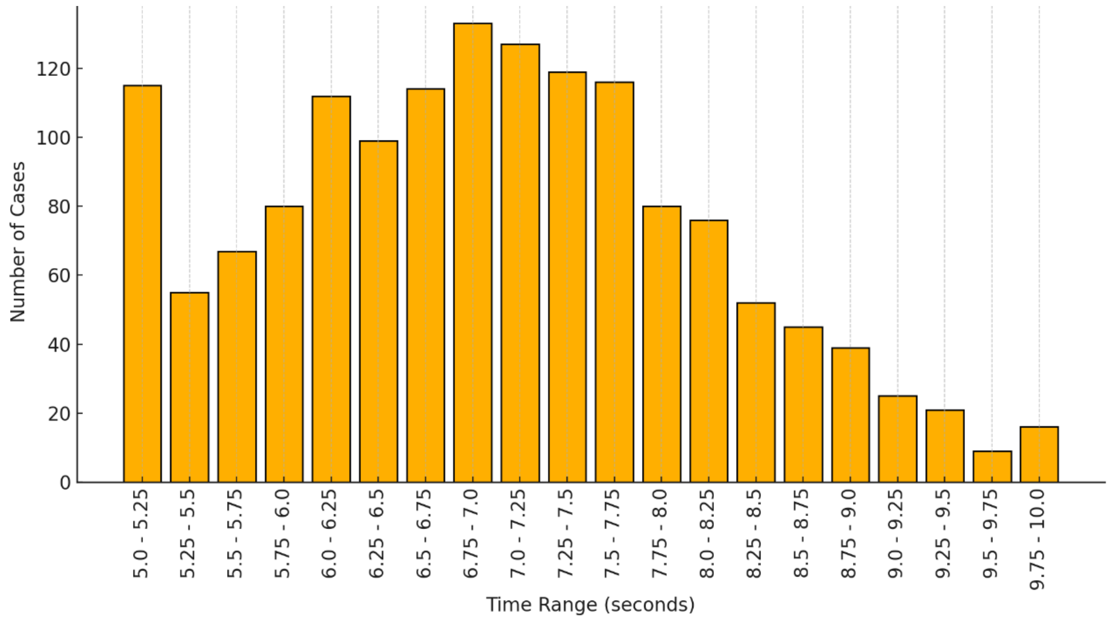

Figure 2: Distribution of Time-to-Diagnosis for AI Chest X-ray Interpretation

Figure 2 shows how often AI inference times were recorded for 1,500 chest X-ray cases. Most predictions are between 6 and 8 seconds, and the average time is 6.9 ± 1.2 seconds. The results are very similar across all patients, meaning the system is well-suited for fast clinical triage in urgent or limited-resource situations.

User Experience and Adoption

All 23 participants (100% of them) provided structured feedback on the survey, and these included 20 healthcare professionals and 3 IT staff. Users gave the AI interface an average score of 4.4 for ease of use, 4.1 for trust in its predictions, and 4.6 for their willingness to adopt it (Table 3: User Survey Summary). Many users pointed out that the system helped reduce delays in diagnosing patients, especially when large numbers of patients were seen. While the survey findings are positive, they are based on a limited and non-random group, and conducting more user studies at different sites could help apply these results more widely. Figure 3 visualizes user sentiment across the survey categories. Table 3 shows the feedback from 23 users (20 healthcare providers and 3 IT staff) on their experience with the AI system. The scores and standard deviations are provided to help you use the results, trust the AI, and decide if you want to use the tool in your practice.

Table 3. User Survey Summary on AI System Usability and Acceptance

Survey Item	Mean Score (1–5)	Standard Deviation
Ease of Use	4.4	0.5
Trust in AI	4.1	0.6
Willingness to Adopt	4.6	0.4

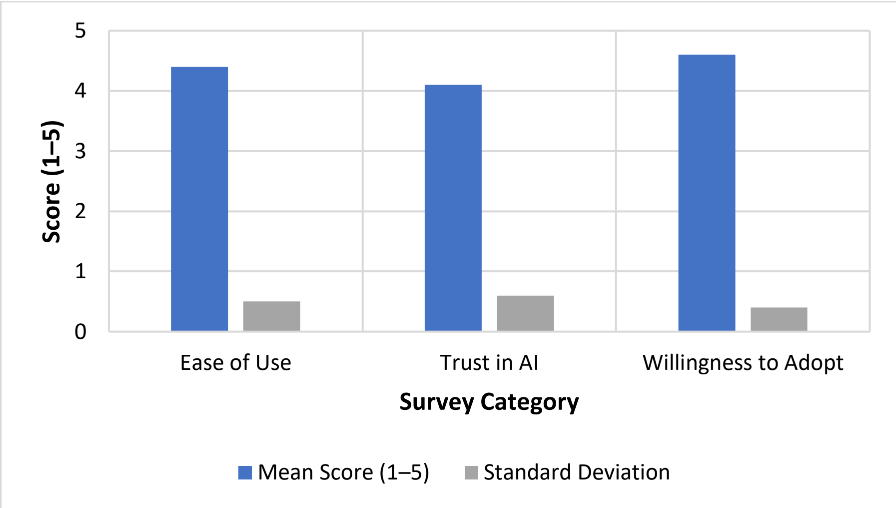

Figure 3: User Feedback on AI System Usability and Acceptance

Figure 3 shows a bar chart that compares the mean scores and standard deviations for the three categories users rated: ease of use, trust in AI predictions, and willingness to adopt the system. 23 respondents were asked to rate their experiences on a 5-point Likert scale. The scores are high because users are happy and ready to use the product.

System Stability and Scalability Observations

The system continued to work smoothly during the deployment without needing any changes or updates. All five sites operated autonomously with minimal technical support. The training session was only 1 hour long, and all technicians were able to use the system after that. With offline-first architecture, systems could work without constant internet access, which made them suitable for expansion in other similar health systems.

Discussion

It has been shown that a low-cost, offline-functioning AI system is suitable for chest X-ray interpretation in rural healthcare centers that lack internet access.[9] The AI model was able to diagnose accurately, with a sensitivity of 88%, a specificity of 85%, and an AUC of 0.91. It took the system only seven seconds to process each image, kept five sites running at 97% uptime, and was praised by both healthcare staff and IT workers. The results back up our hypothesis that an AI-based diagnostic tool can work effectively and efficiently where there are no radiologists and regular internet access. The outcomes highlight that this approach could be used in many regions to help more people get access to diagnostics.[10,11]

To appreciate what these outcomes mean, we look at similar studies in AI radiology and consider how they affect global health. Our results are consistent with previous studies in AI-based chest X-ray analysis, especially those using CheXNet[7], which also achieved high accuracy on standard datasets.[12] Most studies done before have looked at past events in urban hospitals that have advanced digital systems.[13,14] On the other hand, our study checked the AI system in real-life settings at low-resource clinics using inexpensive, offline equipment.[15] Using this approach, the literature gains a strong implementation science perspective.[16],[17] It highlights that diagnostic AI can be useful in places where standard radiology assistance is lacking, which is not well explored in current studies.[18]

AI helps frontline health workers notice abnormalities, speed up triage, and make sure patients are referred on time.[19] It can be especially useful for finding tuberculosis, pneumonia, and trauma early, which are common in rural healthcare.[20,21] This work supports WHO's efforts to make essential diagnostics more accessible and helps achieve Sustainable Development Goal 3, which aims for universal health coverage.[22] The fact that the tool is inexpensive, can work without cloud infrastructure, and is compatible with open-source components makes it a good choice for use in national screening programs and NGO-led projects in low- and middle-income countries (LMICs).[23]

The system was used as part of regular clinical work to assess how it worked and how reliable it was. This distinguishes it from controlled lab-based studies.[24] The study also looked at how easy the tool was to use and found that frontline staff were very willing to trust and use it. The project was built using open-source software and budget-friendly hardware, which made it easy for others to reproduce and use.[25]

Despite its promising outcomes, the study has limitations. Even though the test set was enough for the first evaluation, it was small and consisted of data from only one regional archive, which limited how widely it could be used. Only chest X-rays were considered in the analysis, since there are no validated models for ultrasound, which is commonly used in maternal and emergency care.[26] The AI system was not checked for accuracy in situations where the images are low-quality or show children. But the fact that the model works well everywhere and users are engaged means it can be tested in different settings. More research should be done in various locations to see how this system works for different groups of people. More research should be done using ultrasound and other imaging methods to improve the clinical value of these techniques. It is important to evaluate over time how AI-based triage influences the start of treatment, patient results, and the efficiency of the health system. Using federated learning could allow models to adapt better and still protect sensitive data. Partnering with national health authorities will be important to add these systems to digital health strategies, electronic medical records, and regular care pathways. The research proves that a low-cost AI radiology solution is both possible, accurate, and practical in real-life settings where resources are scarce. Because it is easy to use, fast, and requires little training, the system could help increase medical diagnosis in areas with limited resources. When these tools are further confirmed and supported by policies, they could assist in addressing the global shortage of diagnostic services and promoting equal healthcare.

Conclusion

This research tested the hypothesis that an offline, low-cost AI radiology system could be used to augment diagnostic decision-making in environments without radiologists or reliable internet. Here, the authors describe a real-world deployment of such a system in five rural health centers, where the AI model proved highly accurate, fast in inference, and reliable in difficult conditions. Frontline clinicians, with limited training, highly rated the system for usability and integration into standard clinical workflows. In contrast to prior studies limited to retrospective analysis or highly digitized settings, this work uniquely demonstrates the feasibility of AI integration in low-resource settings. Developed on open-source software and low-cost edge hardware, the system is independent of cloud infrastructure, recommending it for use in primary care in low- and middle-income countries. Its high user acceptance and operational stability support its promise for sustainable scale-up. By closing the gap between algorithmic innovation and real-world implementation, this research provides a reproducible model for AI-facilitated diagnostics aligned with universal health equity objectives and the WHO priority. It not only confirms a scalable approach to radiological interpretation for use in underserved environments but also sets the stage for potential extension to other imaging modalities and health system environments.

References

1. Wong E, Bermudez-Cañete A, Campbell MJ, Rhew DC. Bridging the Digital Divide: A Practical Roadmap for Deploying Medical Artificial Intelligence Technologies in Low-Resource Settings. Popul Health Manag. 2025.
2. Khan FS, ul Hassan M, Ahmed A, Khan AA. Bridging the Gap Solutions for Emergency Care in Resource-Limited Health Care Environments. Rev J Neurol Med Sci Rev. 2025;3(1):253–62.
3. Lamichhane B, Neupane N. Improved healthcare access in low-resource regions: A review of technological solutions. arXiv Prepr arXiv. 2022;2205.10913.
4. Xiques-Molina W, Lozada-Martinez ID, Fiorillo-Moreno O, Hernández-Lastra AL, Bermúdez V. Operational Advantages of Novel Strategies Supported by Portability and Artificial Intelligence for Breast Cancer Screening in Low-Resource Rural Areas: Opportunities to Address Health Inequities and Vulnerability. Medicina. 2025;61(2):242.
5. Valerio JE, Ramirez-Velandia F, Fernandez-Gomez MP, Rea NS, Alvarez-Pinzon AM. Bridging the global technology gap in neurosurgery: disparities in access to advanced tools for brain tumor resection. Neurosurg Pract. 2024;5(2):e00090.
6. Ralph-Okhiria O, Alonge I. Leveraging artificial intelligence to strengthen surgical systems in sub-Saharan Africa. Acad Med. 2025;2(2).
7. Venkatayogi N, Gupta M, Gupta A, Nallaparaju S, Cheemalamarri N, Gilari K, et al. From seeing to knowing with artificial intelligence: a scoping review of point-of-care ultrasound in low-resource settings. Appl Sci. 2023;13(14):8427.
8. Kim S, Fischetti C, Guy M, Hsu E, Fox J, Young SD. Artificial intelligence (AI) applications for point of care ultrasound (POCUS) in low-resource settings: a scoping review. Diagnostics. 2024;14(15):1669.
9. Krones F, Walker B. From theoretical models to practical deployment: A perspective and case study of opportunities and challenges in AI-driven healthcare research for low-income settings. medRxiv. 2023.
10. Ohene-Botwe B, Amedu C, Antwi WK, Abdul-Razak W, Kyei KA, Arkoh S, et al. Promoting sustainability activities in clinical radiography practice and education in resource-limited countries: a discussion paper. Radiography. 2024;30:56 61.
11. Pillay TS, Khan A, Yenice S. Artificial intelligence (AI) in point-of-care testing. Clin Chim Acta. 2025;120341.
12. Were MC. Challenges in digital medicine applications in under-resourced settings. Nat Commun. 2022;13:3020.
13. Kiyasseh D, Zhu T, Clifton D. The promise of clinical decision support systems targeting low-resource settings. IEEE Rev Biomed Eng. 2020;15:354 71.
14. Reddy KJ. Accessibility and Affordability of Technologies. In: Innovations in Neurocognitive Rehabilitation: Harnessing Technology for Effective Therapy. Cham: Springer Nature Switzerland; 2025. p. 285–304.
15. Dave D, Sawhney G. Revolutionizing Diabetic Retinopathy Diagnosis in Third World Countries: The Transformative Potential of Smartphone-Based AI. Authorea Preprints. 2023.
16. Vohra H, Hasan MK, Abdullah SNHS, Islam S, Abd Rahman AH, Bhojake SN, et al. A low-cost AI-Powered System for Early Detection of Diabetic Retinopathy and Ocular Diseases in Resource-Limited Settings. IEEE Access. 2025.
17. Opia FN, Igboekulie CV, Matthew KA. Socioeconomic Disparities in Breast Cancer Care: Addressing Global Challenges in Oncology Outcomes.

18. Alabdaljabar MS, Hasan B, Noseworthy PA, Maalouf JF, Ammash NM, Hashmi SK. Machine learning in cardiology: a potential real-world solution in low- and middle-income countries. J Multidiscip Healthc. 2023;285–95.
19. Opia FN, Matthew KA, Matthew TF. Leveraging algorithmic and machine learning technologies for breast cancer management in Sub-Saharan Africa. Int J Multidiscip Compr Res. 2022.
20. Grover S, Court L, Amoo-Mitchell S, Longo J, Rodin D, Scott AA, et al. Global Workforce and Access: Demand, Education, Quality. Semin Radiat Oncol. 2024 Oct;34(4):477–93.
21. Almaghadi FA, Almajhadi KAS, Alsukaibi KF, Alsalman AHA, Al Haidar MHH, Alzarraj YI, et al. The Invisible Crisis: Overcoming Technological and Systemic Hurdles in Radiology and Optometry. J Int Crisis Risk Commun Res. 2024;7(S3):736.
22. Agbeyangi A, Suleman H. Advances and Challenges in Low-Resource-Environment Software Systems: A Survey. Informatics. 2024 Nov;11(4):90.
23. Ranjit S, Kissoon N. Challenges and Solutions in translating sepsis guidelines into practice in resource-limited settings. Transl Pediatr. 2021;10(10):2646.
24. Anazodo U, Asllani I, Fatade A. Building a Global Power of Experience in Diagnostic Imaging–Lessons from Africa’s COVID-19 Response.
25. Nhat PTH. Translation of Computer-Assisted Point-of-care Ultrasound Imaging Methods in a Resource-Limited Intensive Care Unit.
26. Alaswad M, Abady EMA, Darawish SM, Barabrah AM, Okesanya OJ, Zehra SA, et al. Point-of-Care Ultrasound in Pediatric Emergency Care in Low-Resource Settings: A Literature Review of Applications, Successes, and Future Directions. Curr Treat Options Pediatr. 2025;11(1):1–13.