Patterns

CellPress
OPEN ACCESS

Perspective

Optimizing human-centered AI for healthcare in the Global South

Chinasa T. Okolo[1,*]
[1]Department of Computer Science, Cornell University, Ithaca, NY 14853, USA
*Correspondence: chinasa@cs.cornell.edu
https://doi.org/10.1016/j.patter.2021.100421

THE BIGGER PICTURE Despite growing enthusiasm to address societal problems using AI, there is a scarce amount of research studying the implications and challenges associated with integrating AI-enabled technologies into low-resource communities throughout the Global South. Neglecting to analyze the unique needs and requirements of the frontline workers expected to operate AI systems, especially those used for healthcare, stands to exacerbate existing issues in algorithmic bias and impose additional work burdens, deteriorating the level of care provided to vulnerable communities.

1 2 3 4 5 **Concept:** Basic principles of a new data science output observed and reported

SUMMARY

Over the past 60 years, artificial intelligence (AI) has made significant progress, but most of its benefits have failed to make a significant impact within the Global South. Current practices that have led to biased systems will prevent AI from being actualized unless significant efforts are made to change them. As technical advances in AI and an interest in solving new problems lead researchers and tech companies to develop AI applications that target the health of marginalized communities, it is crucially important to study how AI can be used to empower those on the front lines in the Global South and how these tools can be optimally designed for marginalized communities. This perspective examines the landscape of AI for healthcare in the Global South and the evaluations of such systems and provides tangible recommendations for AI practitioners and human-centered researchers to incorporate in the development of AI systems for use with marginalized populations.

INTRODUCTION

As artificial intelligence (AI) and machine learning (ML) continue to spread widely into domains, such as agriculture, government, and healthcare, most conversations regarding its implications focus on communities in the Global North, such as the US and Europe. Nonprofits, tech companies, and governments are rushing to build and deploy AI systems, yet they fail to examine the knowledge, needs, and perceptions of the frontline workers who will be expected to operate these systems. This approach risks imposing extra work and inefficiencies at point of care or causing harm to the very communities these technologies intend to serve. Thus, as technical advances in AI and an interest in solving new problems lead researchers and tech companies to develop AI applications that target the health of marginalized communities, it is crucially important to study how AI can be used to empower those on the frontlines in the Global South and how these tools can be optimally designed for marginalized communities. Efforts toward these problems have primarily been driven by large tech companies, such as Google, IBM, and Microsoft, further hindering the ability of researchers from marginalized backgrounds to take agency over AI development affecting their communities. This emphasizes the need of local stakeholders to be centered and prioritized in the process of building technical solutions, inverting existing practices of AI development.

This article highlights human-centered work in the intersection of AI and healthcare to explore the infrastructural, cultural, and technological factors that impact the delivery of AI-enabled services in the Global South. While synthesizing existing literature, it provides actionable steps for both AI practitioners and human-centered researchers to reflect on and implement in their respective solutions addressing healthcare issues within low-resource environments. With a focus on centering the end users of AI-enabled systems for healthcare, the solutions provided in this paper are primarily non-technical but will have a significant impact in shaping the technical aspects of such systems. The recommendations provided within this article leverage insights presented from previous work, providing a deeper perspective on the implications of deploying AI systems for healthcare in the Global South.

The contribution of this work lies in its calls to increase cross-disciplinary collaboration in the conduct of AI research, implement design considerations that provide users with ways to maintain their autonomy in automated decision making, localize practices of AI development by not replicating habits seen in the development of systems in the West, consider the utility of creating AI to complement rather than replace human work, and invest adequately in frontline workers to equip them with the knowledge necessary to operate AI-enabled tools and manage concerns associated with them. We begin this article by providing an overview of AI-for-health in the Global South and move on to provide insight into human-centered evaluations conducted within this field and issues regarding the lack of evaluations. Next, we discuss how to bring awareness to known harms of AI research in the Global South and ways to mitigate such issues by encouraging participatory design, finalized with a conclusion envisioning AI in human-centered healthcare.

BACKGROUND

AI in healthcare in the Global South

The realization of AI in the Global South has led to a variety of use cases throughout agriculture, education, financial services, and, most notably, healthcare. Within this region, AI has been leveraged to support maternal and child health,[1] engagement with health programs,[2] and for the ongoing COVID-19 pandemic.[3–6] There has also been a significant focus on diseases that disproportionately affect communities in the Global South, such as dengue fever,[7] ebola,[8] tuberculosis,[9–12] and diabetic retinopathy,[13] a disease that causes damage to the blood vessels as a complication of diabetes. To address this issue, Bellemo et al. developed a model consisting of a convolutional neural network to identify cases using retinal image scans.[14] This tool was validated in a study conducted with the Community Eye Service Programme in Zambia, demonstrating similar diagnostic capabilities to those of human graders. Other work by Gulshan et al. has focused on screening diabetic patients in India and showed comparable results with the study conducted in Zambia, with their respective deep learning models showing accuracy equal to or exceeding that of human graders.[15] With this growth in efforts targeting diseases relevant to this region, there is hope that future AI tools will target neglected tropical diseases that also disproportionately affect the Global South, such as leishmaniasis, schistosomiasis, and mycetoma.

A growing field of literature has begun to focus on community health workers; the frontline workers who will be most responsible for using and integrating AI-enabled tools into healthcare services in the Global South. Ismail and Kumar outline the use of AI in frontline health throughout the Global South, analyzing the motivations behind such work, the stakeholders involved, and how well such applications engage with local communities.[16] The authors base this research on their extensive background conducting ethnographic fieldwork with community health workers in India, methods not commonly used by AI researchers. Their paper also provides actionable design considerations for AI systems being developed for social good that expand beyond healthcare, setting the stage for more impactful AI applications integrated in low-resource regions and with disadvantaged populations. Okolo et al. conducted an ethnographic study in rural India, characterizing the knowledge, perceptions, and understanding that community health workers hold about AI and its potential benefits and drawbacks.[17] That article was the first of its kind to directly examine community health workers in this manner, providing valuable insights for not only designers and developers of AI systems but for the governments and policy makers who will govern the use of such systems.

Within the Global South, non-governmental organizations (NGOs) and private foundations play a significant role in funding and organizing the delivery of healthcare services. These stakeholders have also realized the potential of AI within healthcare and have begun to actively work toward understanding the best way to develop technical solutions sustainably. A report created in partnership with The Rockefeller Foundation, the United States Agency for International Development's, and the Bill and Melinda Gates Foundation, surveys current practices of AI in healthcare, examining the critical challenges associated with scaling AI development in the Global South and the investment needed to effectively support such work.[18] Other work by Williams et al. has surveyed deep learning techniques for medical image analysis that democratize access to AI in low-resource regions, allowing those with less technical expertise to successfully navigate these tools.[19] Williams et al. also focus on the role that NGOs can play in identifying problems to which AI is a possible solution. With the shift to actualizing human-centered ML in healthcare, work centering community health workers and other frontline healthcare professionals will be essential in moving the field forward and ensuring that AI interventions work equitably across the Globe.

Evaluations of healthcare AI systems

As AI continues to make the transition from the lab to the real world, its impact and accuracy have become increasingly scrutinized. However, human-centered evaluations, analyses that work to understand how humans use and are impacted by these technical systems in real-life settings, have been less common, especially in healthcare. The first known study to observe an AI system in clinical use was published in 2020.[20] Compared with the thousands of ML models that have been developed for medical contexts, there are significantly fewer published studies that have observed AI systems in deployment with patients and healthcare professionals and even fewer that have done this in the context of the Global South. Motivated by issues with screening for diabetic retinopathy in Southeast Asia, researchers at Google AI applied deep learning techniques to shorten the often weeks-long process of having an ophthalmologist review retinal scans.[15] With the algorithm showing over 90% sensitivity and specificity, it was perceived to be comparable with a physician, initiating a clinical evaluation. In partnership with 11 clinics in Thailand, researchers from Google Health conducted a human-centered observational study to analyze how the introduction of this algorithm impacted clinical workflows and the factors that affected algorithmic performance of the system.[20] The authors of this study claim that it is the "first human-centered observational study of a deep learning system deployed directly in clinical care with patients," providing significant findings for both AI-for-health research and AI research in the Global South.

While the system developed by these Google researchers was thought to be beneficial from preliminary results seen in the lab, it could not be described as "useful" until it had been tested in clinics, working with real patient data. Despite the researchers performing fieldwork pre- and post-deployment of the algorithm, there were significant issues that occurred in the integration of the system. A major finding from this study showed that the system poorly handled low-quality retinal images taken in the clinics, as the underlying algorithm had originally been trained on high-quality lab images. The inability of the system to handle images produced from the clinic led to major disruptions during and after the screening process, affecting both patients and nurses. For example, poor lighting conditions often led to low-quality images, increasing the time nurses had to spend re-taking images and increasing the overall waiting time for other patients. Subsequently, the flash from re-taking photos led to discomfort for patients, which was recognized by the nurses who limited such attempts to two tries. Additional factors such as internet connectivity, capability of equipment to capture sufficient images, and availability of medical professionals to confirm results from the algorithm also impacted the success of this AI system.

Moreso, the study from Google AI/Health highlights that, in addition to the accuracy of ML models, the ability to improve patient care is a significant factor in determining the success of a deep learning model. More specifically: "[w]hen deploying deep learning systems, end-users and their environment determine how a new system will be implemented; that implementation is of equal importance to the accuracy of the algorithm itself."

This statement by the authors is not only useful for the broader AI for healthcare research field, but especially for AI tools being developed for and in the Global South. End users and their respective environments are under-studied in the deployment of deep learning systems, and as these deployments into the "real world" increase, more effort should be taken by both human-computer interaction (HCI) and AI/ML researchers to evaluate the models in context. In addition, the importance of incorporating methods from the social sciences, and working with cross-disciplinary researchers, such as sociologists and anthropologists, should be encouraged and embraced by technologists. These researchers bring expertise in methods, such as ethnography and field surveying, that would be useful for analyzing the sociotechnical impacts of AI systems, improving their integration into local communities.

TESTING AI-ENABLED INTERVENTIONS

Need for evaluation

A primary issue with current AI development is the significant lack of understanding about how AI works in the real world and a further lack of understanding about how AI works in the Global South. With AI applications increasingly adopted for use in low-resource environments, it becomes even more important to test and evaluate these technologies with the communities and in the environments in which they are expected to be integrated. For healthcare concerns, this need becomes even greater due to the critical nature of medical decision making. Even in what may seem to be the most optimal or beneficial case, it is particularly difficult to integrate AI tools for patient care. A commentary by Hu et al. explains the shortage of AI models developed for COVID-19 diagnosis in frontline healthcare services, highlighting the lack of transparency in the data sources used to train such models and the inflexibility of approaches that are not robust for a wide variety of situations.[21] In a paper summarizing the challenges of introducing AI systems into healthcare, Kelly et al. note the lack of prospective studies that have been conducted with AI systems incorporating real-world data into their respective analyses.[22] This current focus on using established benchmark medical datasets for retrospective studies stands to continue producing ML models that are unable to adapt to situations not represented in training data and to non-ideal clinical settings. While AI development and testing have primarily been conducted in the Global North, researchers in low-resource countries have begun to leverage the power of AI over the past few years. Work done by researchers in Brazil developed 22 AI models for healthcare and tested them in real-world settings for an average of 2 years.[23] This in-depth evaluation analyzed the difficulties of using AI at scale in healthcare and provided a set of tangible requirements for scaling AI for such use. The proposed requirements: consistency, reproducibility, versioning and traceability, preservation of data quality, scalability, and antifragility are bolstered by the fact that they came out of the deployment of AI in practical healthcare contexts, setting a model for future AI-for-health studies to follow.

Evaluation issues

In cases where AI is evaluated in medical contexts, previous work has found significant issues of bias[24] and irreproducibility.[25] This raises serious concerns about the state of ML research in healthcare, and researchers should be wary of deploying AI-enabled technologies in environments that have limited regulatory oversight. AI systems for healthcare continue to be commercialized and, without rigorous oversight, could result in harm for the communities at the point of care. An independent investigation of the Epic Sepsis Model, an early-warning system for sepsis that has been widely implemented across hospitals in the US, highlights the need for rigorous evaluation of AI models used in healthcare and the need to center the end users in the design and development of such systems.[26] While it is extremely concerning that a system with such poor performance (AUC = 0.63) would be integrated for use in a real-world setting, the impact of the constant false alarms that overwhelmed physicians is worth noting. This stresses the need for AI practitioners and user design (UX) researchers to consider the impact of AI systems on end users and create space for these users to leverage their autonomy in the control of such systems.

The democratization of AI development has led to solutions being quickly developed for a myriad of real-life use cases. The acceleration of the COVID-19 pandemic led to a barrage of AI models being proposed for use in aiding diagnosis and treatment. While medical journals and preprint servers prioritized the publication of such studies, work by researchers at Cambridge University analyzing over 400 models developed to diagnose COVID-19 showed that every single one was not suitable for clinical use.[27] Given the time frame and lack of validation experiments for many of these studies, it is unlikely that they would have been deployed in real-world systems. With the gravity of the current pandemic and the impact it has had on countries in the Global South, such as Brazil, India, and South Africa, there is a possibility that some of these flawed models may have been considered to combat the enormity of cases. Many issues with evaluating AI models for healthcare are compounded by the fact that the data used to train these models come from a limited population. In the US, the vast majority of AI diagnostic tools rely on data from three states: California, Massachusetts, and New York.[28] While these states make up a significant percentage of the US population, there are a variety of environmental and socioeconomic factors that impact patients outside of these states that could potentially lead to negative outcomes on those not represented in the training datasets. As AI development moves from the Global North to the Global South, if such practices are copied, the potential for negative outcomes could be even greater due to the vast differences in infrastructure, climate, social norms, and resources available to people living in these regions. With this in mind, it becomes even more imperative to prioritize the local contexts of where AI tools are integrated.

DISCUSSION

There is much room to reshape how AI developers and healthcare practitioners approach current development of AI solutions for healthcare. AI researchers and the AI research community at large hold a significant role in shaping research agendas aligned with current initiatives of AI for Social Good and should do more to mitigate potential harms associated with these initiatives.

Understanding the known harms and impact of AI-for-health interventions

As the ethical implications of AI systems continue to be prioritized in AI development, albeit mainly in Western contexts, it is important that this work also be highlighted in AI for healthcare. Recent work by Gichoya et al. has called out the need for regulatory guidelines and protocols governing the use of AI in healthcare to engage more closely with issues of bias, fairness, and disparate impact.[29] It is also extremely important for AI practitioners working on problems centered in the Global South to understand and work proactively toward ameliorating issues of fairness, accountability, transparency, and ethics in AI. In low-resource settings where values, such as inclusion and privacy, may not align with Western norms, research that anticipates the risks associated with these differences will become essential to ensuring fairness in such work. Research done by Sambasivan et al. show that typical Western takes on algorithmic fairness are not portable in the context of the cultural values and standards in India.[30] Similar conclusions may hold true for other regions within the Global South. In regions where ethnic, tribal, or religious affiliation hold more power in society compared with Western notions of race, which do not apply in countries such as Nigeria or Pakistan, these demographics should be seriously considered in the development, deployment, and use of AI systems.

Existing work in AI for Social Good brings up questions challenging the role of responsibility and accountability in AI research and outlining what this looks like in non-Western contexts. Issues such as infrastructural limitations, community harm, power asymmetries, and more must be identified and examined. Over the past decade, community health workers have increasingly adopted mobile phones in their work, but this penetration of mobile phone use has been limited to regions where there is sufficient funding for such devices and where there is adequate infrastructural capacity. When developing AI-enabled interventions for use in low-resource contexts, the needs of people operating these tools should be centered, especially if they are novice technology users. Sufficient investments into adequate training and compensation for community health workers can avoid imbalanced incentive structures that have traditionally favored researchers over end users. In the case of older community health workers who have to learn how to operate new devices and the sophisticated algorithms associated with them, younger workers may take on the role of IT support while also having to manage their own respective caseloads. When AI systems are transported from high- to low-resource domains, infrastructural limitations can lead to incorrect predictions and system failures that complicate workflows and cause stress to both healthcare workers and patients. Misdiagnoses not only cause emotional distress but can lead to financial distress if patients have to spend time away from work traveling to healthcare centers for further diagnostic testing.

Understanding how to account for trust in AI systems used in healthcare is an additional area of exploration that should be pursued further by the AI research community. While some research has shown that medical advice generated by AI models is generally labeled as lower quality by radiologists,[31] research conducted in the Global South has shown the potentially negative compounding effects that could occur due to the limited knowledge community health workers possess about AI as a technology.[17] This may lead community health workers and other healthcare professionals to lend unwarranted authority toward decisions produced from AI systems. Even more worrying is that incorrect decisions produced from such systems may be overlooked or left unchallenged if these workers doubt their own abilities in the presence of systems they perceive to be highly advanced. Overall, when thinking of these issues, the developers and designers of AI systems should build them in such a way that these systems adapt to the contexts in which they will be used, and ensure that users learn how to use them effectively and in a way that leverages their autonomy.

As efforts grow for institutions to develop more inclusive ML models, there remains a lack of focus on representation from regions outside the Global North. One of the most important parts of ML development: data collection and processing, is inequitable and often inaccessible to researchers in the Global South. In their work interviewing AI practitioners in Southeast Asia and Sub-Saharan Africa, Sambasivan et al. highlight the effects of data inequity when applying commonplace ML practices in the Global South.[32] The negative impacts of such systems stand to threaten the potential for AI to be effective in high-stakes situations, such as healthcare in low-resource regions. Work by Mitchell et al. emphasizes the need for ML models, especially those trained on data from a single source, to be validated on local, representative datasets.[33] This will pave the way for generalizability in ML applications for healthcare, while ensuring that the Global South is equitably represented in these solutions.

Amid accelerated data collection efforts for ML in the Global South, the implications of this work need to be accounted for. The lack of geographical diversity in ML datasets for healthcare is concerning and should be rectified. A review of datasets used for ophthalmological imaging describes this situation as a form of "data poverty," where countries or regions are underrepresented in datasets.[34] The effects of data poverty have the potential to lead to ML models solely focusing on diseases that affect Western, richer countries, worsening global health equity. Recent work has directly interfaced with AI practitioners to understand data practices on the African continent and the accompanying power imbalances between researchers and funding bodies that limit how such data can be used to drive solutions to local problems.[35] That work also highlights how Western narratives dominate data-sharing practices, often discarding local context and expertise, a recurring yet concerning trend. In regions such as India, frontline healthcare workers, many of whom are women, lead efforts for collecting health data. Ismail and Kumar expose the challenges these workers face, providing valuable insights into their data collection practices, their relationships with other workers, and interactions with local residents.[36] Issues such as language barriers and healthcare literacy affected both the accuracy and completeness of such data, identifying additional areas for AI practitioners to focus on. The end users of AI systems should be given agency not only over the models themselves, but the resulting data produced from them. AI researchers and developers also have the responsibility to make concerns regarding privacy and security of sensitive data aware to users and provide communities with actionable ways to manage these concerns.

Encouraging participatory design

Many of the issues seen in the implementation of AI systems could possibly be alleviated by introducing practices of participatory design. Participatory design of intelligent systems should be a forethought and should work to actively include all stakeholders at each step. As the field of AI moves toward being "human-centered," there stands to be much benefit if these principles are taken more seriously. Previous work by Bondi et al. examines work in AI for Social Good projects, proposing a framework to guide the establishment of AI projects within low-resource environments that prioritize the needs of, and enhances the capabilities of, local communities.[37] In healthcare, new methods are being introduced to improve integration of AI-enabled tools, which often fail in real-life clinical settings. Methods by Jacobs et al. focus on using an iterative co-design process to aid the design of tools to support clinical decision making.[38] Similar methods may prove useful in situations in the Global South where there is a significant lack of qualified medical professionals and the needs of the frontline workers who serve communities are not well understood. Co-designing AI tools in these situations can help ensure that they are effective in a variety of contexts.

While participatory design is helpful in improving the adoption of technical solutions, there are additional limitations, such as the potential for exploitation, lack of short-term gain to local communities, and an increased burden on marginalized populations, that must be considered. Pierre et al. utilize the concept of epistemic burden to describe the intensive process of collecting, curating, and disseminating data in participatory design projects.[39] Their work reflects on previous research projects conducted by the authors and incorporates case studies to highlight the extra burden inadvertently placed on community partners, providing insights to evade potentially negative impacts. Other work highlights the power imbalances that continue to exist in participatory design research. Harrington et al. posit that current practices of participatory design primarily cater to those with significant privilege and financial resources, often failing to deliver equitable solutions for underprivileged populations.[40] In the context of the Global South, the findings from these papers provide a strong framework to understand the issues that marginalized communities experience in the participatory design process and offer proactive measures for researchers working in this domain to incorporate.

AI is often introduced with the mindset of replacing human work rather than complementing it. While automated systems may be useful for dangerous work such as manufacturing or mining, in human-centered fields such as healthcare and education, full automation may have a negative net impact. In research describing AI systems, there is strong potential for high efficacy of human-AI systems working together rather than alone.[41,42] In research conducted with community health workers in rural India, results from the study showed that these workers embrace the idea of working together with AI-enabled applications, acknowledging that such systems would be unable to replace their work, such as consultations, immunizations, and discussions on sensitive topics, including breastfeeding and family planning.[17] These findings possibly indicate that in low-resource environments throughout the Global South, introducing fully autonomous AI systems may harm patient-provider relationships and introduce the possibility for degraded levels of care. Shifting from an ethos of replacement to complementation in the development of AI systems paves the way for participatory design to take precedence.

Data scientists and AI/ML practitioners should be encouraged to collaborate directly with the communities their work intends to impact. Collaboration should not be an afterthought and needs to be actively implemented throughout the entire design, development, and integration process. Participatory design also necessitates the formation of interdisciplinary teams. To prevent AI work from being conducted in silos and with limited perspectives, it is necessary that AI researchers initiate collaborations with social scientists and HCI researchers. These efforts will help to increase the number of human-centered evaluations of AI systems developed for medical use in low-resource contexts and provide a more comprehensive understanding of the potential benefits and implications of AI in these regions. Upstream development challenges, such as infrastructural capacity, technical and AI literacy of end users, and even the necessity or feasibility of AI as a solution for certain healthcare problems, should be taken into consideration by data science researchers. While AI has shown the potential to make tremendous contributions to healthcare and medicine, most of this work has not been actualized due to ML models being trained on unrepresentative data and evaluated in lab settings that are very different from real-world deployment contexts. To arrive at a place where AI can be truly leveraged "for good" in these domains, it is imperative that human-centered methods be continually developed to adapt to this ever-changing landscape.

A VISION FOR AI IN HUMAN-CENTERED HEALTHCARE

While significant progress in AI has been made in domains ranging from precision agriculture to structural biology, most of its benefits have not reached the Global South and will not be actualized unless efforts are made to change current practices. At the present moment, the challenges seen in current AI practices, especially those regarding bias, stand to exacerbate existing inequities in low-resource regions. Centering humans in the development of artificially intelligent systems, specifically those intended for healthcare, shows promise for improving current systems but will not be useful unless the distinct needs of communities are met. Leveraging these methods for AI in healthcare has proven to be challenging, as demonstrated in the existing literature. These challenges are further compounded by the infrastructural and knowledge gaps present in the Global South that currently hamper effective AI integration. However, the increase in research specifically centering frontline workers in the Global South gives AI researchers, HCI scholars, and UX designers new insights to inform the development of AI-enabled applications.

As AI becomes integrated within healthcare systems in the Global South, frontline healthcare workers will be expected to operate AI-enabled tools. However, their use of such systems will not be effective unless they are able to independently operate and comprehend the decisions produced by such tools. As designers and developers build these systems, it is imperative that the needs, values, and concerns of these users are addressed. Participatory design methods show promise to address concerns raised by researchers who have interacted with novice technology users, but the expense associated with human-centered research along with the cost of implementing systems may inhibit the actualization of these goals. This perspective employs previous research in AI for healthcare, human-centered AI, and participatory design to understand the sociotechnical factors that impact successful implementation of AI systems in the Global South and to provide recommendations for AI and HCI practitioners to enable the capacity of frontline workers.

Structural changes in the incentives that reward model-driven research within the AI community can help shift AI research toward an agenda that embraces the integration of human-centered research within the model development pipeline. The introduction of broader impact statement requirements at premier conference venues, such as the Conference on Neural Information Processing Systems (NeurIPS), and special tracks on "AI for social impact" at the Association for the Advancement of Artificial Intelligence Conference on Artificial Intelligence (AAAI) and the International Joint Conference on Artificial Intelligence (IJCAI) are initial steps toward encouraging reflection during the AI development process and bringing systems closer to actuality. However, the implementation of AI models and systems in real-world settings is still under-explored. Tracks or reviewing criteria at these conferences could be introduced to reward constructive implementations that center on the needs of local communities. These efforts would motivate researchers toward addressing the challenges laid out within this perspective, advancing the role of AI and how it could potentially translate to improved healthcare services within the Global South. The suggestions outlined only comprise a minimal amount of the work needed to fully optimize the potential of centering humans in AI for healthcare, but true progress will occur as researchers in AI, HCI, and the social sciences work together toward this agenda.

ACKNOWLEDGMENTS

This paper is dedicated to frontline healthcare workers who work tirelessly on a daily basis to provide essential services to their communities and the AI research communities and organizations within the Global South, navigating barriers to develop solutions that truly meet the needs of their communities. I extend my deepest gratitude to the anonymous reviewers who provided invaluable feedback for this article. Last but not least, I am thankful for the broader AI for Social Good and Human-Computer Interaction for Development (HCI4D) communities that have inspired my respective research agenda and the objectives presented in this manuscript.

REFERENCES

1. Wadhwani, A.I. (2019). Wadhwani AI—maternal, newborn, and child health. https://www.wadhwaniai.org/work/maternal-newborn-child-health/.

2. Nishtala, S., Kamarthi, H., Thakkar, D., Narayanan, D., Grama, A., Hegde, A., Padmanabhan, R., Madhiwalla, N., Chaudhary, S., Ravindran, B., and Tambe, M. (2020). Missed calls, automated calls and health support: using AI to improve maternal health outcomes by increasing program engagement. arXiv, 2006.07590.

3. Moftakhar, L., Mozhgan, S.E.I.F., and Safe, M.S. (2020). Exponentially increasing trend of infected patients with COVID-19 in Iran: a comparison of neural network and ARIMA forecasting models. Iranian J. Public Health *49*, 92–100.

4. Tiwari, S., Kumar, S., and Guleria, K. (2020). Outbreak trends of coronavirus disease—2019 in India: a prediction. Disaster Med. Public Health Preparedness *14*, e33–e38.

5. Yang, Z., Zeng, Z., Wang, K., Wong, S.S., Liang, W., Zanin, M., Liu, P., Cao, X., Gao, Z., Mai, Z., and Liang, J. (2020). Modified SEIR and AI prediction of the epidemics trend of COVID-19 in China under public health interventions. J. Thorac. Dis. *12*, 165.

6. Rasheed, J., Jamil, A., Hameed, A.A., Aftab, U., Aftab, J., Shah, S.A., and Draheim, D. (2020). A Survey on Artificial Intelligence Approaches in Supporting Frontline Workers and Decision Makers for COVID-19 Pandemic (Chaos, Solitons & Fractals), p. 110337.

7. Guo, P., Liu, T., Zhang, Q., Wang, L., Xiao, J., Zhang, Q., Luo, G., Li, Z., He, J., Zhang, Y., and Ma, W. (2017). Developing a dengue forecast model using machine learning: a case study in China. PLoS Negl. Trop. Dis. *11*, e0005973.

8. Buscema, M., Asadi-Zeydabadi, M., Lodwick, W., Nde Nembot, A., Bronstein, A., and Newman, F. (2020). Analysis of the Ebola outbreak in 2014 and 2018 in West Africa and Congo by using artificial adaptive systems. Appl. Artif. Intelligence *34*, 597–617.

9. Aguiar, F.S., Torres, R.C., Pinto, J.V., Kritski, A.L., Seixas, J.M., and Mello, F.C. (2016). Development of two artificial neural network models to support the diagnosis of pulmonary tuberculosis in hospitalized patients in Rio de Janeiro, Brazil. Med. Biol. Eng. Comput. *54*, 1751–1759.

10. Meraj, S.S., Yaakob, R., Azman, A., Rum, S.N.M., and Nazri, A.A. (2019). Artificial intelligence in diagnosing tuberculosis: a review. Int. J. Adv. Sci. Eng. Inf. Technol. *9*, 81–91.

11. Young, C., Barker, S., Ehrlich, R., Kistnasamy, B., and Yassi, A. (2020). Computer-aided detection for tuberculosis and silicosis in chest radiographs of gold miners of South Africa. Int. J. Tuberculosis Lung Dis. *24*, 444–451.

12. Cao, Y., Liu, C., Liu, B., Brunette, M.J., Zhang, N., Sun, T., Zhang, P., Peinado, J., Garavito, E.S., Garcia, L.L., and Curioso, W.H. (2016). Improving tuberculosis diagnostics using deep learning and mobile health technologies among resource-poor and marginalized communities. In 2016 IEEE First International Conference on Connected Health: Applications, Systems and Engineering Technologies (CHASE) (IEEE), pp. 274–281.

13. Natarajan, S., Jain, A., Krishnan, R., Rogye, A., and Sivaprasad, S. (2019). Diagnostic accuracy of community-based diabetic retinopathy screening with an offline artificial intelligence system on a smartphone. JAMA Ophthalmol. *137*, 1182–1188.

14. Bellemo, V., Lim, Z.W., Lim, G., Nguyen, Q.D., Xie, Y., Yip, M.Y., Hamzah, H., Ho, J., Lee, X.Q., Hsu, W., and Lee, M.L. (2019). Artificial intelligence using deep learning to screen for referable and vision-threatening diabetic retinopathy in Africa: a clinical validation study. Lancet Digital Health *1*, e35–e44.

15. Gulshan, V., Rajan, R.P., Widner, K., Wu, D., Wubbels, P., Rhodes, T., Whitehouse, K., Coram, M., Corrado, G., Ramasamy, K., and Raman, R. (2019). Performance of a deep-learning algorithm vs manual grading for detecting diabetic retinopathy in India. JAMA Ophthalmol. *137*, 987–993.

16. Ismail, A., and Kumar, N. (2021). AI in global health: the view from the front lines. In Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, pp. 1–21.

17. Okolo, C.T., Kamath, S., Dell, N., and Vashistha, A. (2021). "It cannot do all of my work": community health worker perceptions of AI-enabled mobile health applications in rural India. In Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, pp. 1–20.

18. United States Agency for International Development (2019). AI in global health: defining a collective path forward. https://www.usaid.gov/cii/ai-in-global-health.

19. Williams, D., Hornung, H., Nadimpalli, A., and Peery, A. (2021). Deep learning and its application for healthcare delivery in low and middle income countries. Front. Artif. Intelligence *4*, 30.

20. Beede, E., Baylor, E., Hersch, F., Iurchenko, A., Wilcox, L., Ruamviboonsuk, P., and Vardoulakis, L.M. (2020). A human-centered evaluation of a deep learning system deployed in clinics for the detection of diabetic retinopathy. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, pp. 1–12.

21. Hu, Y., Jacob, J., Parker, G.J., Hawkes, D.J., Hurst, J.R., and Stoyanov, D. (2020). The challenges of deploying artificial intelligence models in a rapidly evolving pandemic. Nat. Machine Intelligence *2*, 298–300.

22. Kelly, C.J., Karthikesalingam, A., Suleyman, M., Corrado, G., and King, D. (2019). Key challenges for delivering clinical impact with artificial intelligence. BMC Med. *17*, 1–9.

23. Brinhosa, A.F., Fernandes, E.R., de Castro Figueroa, R., Silva, A.V.R., and Drabowski, N.R. (2021). Challenges of applying artificial intelligence responsibly at scale in healthcare. J. Health Inform. *12*, 406–411.

24. Obermeyer, Z., Powers, B., Vogeli, C., and Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. Science *366*, 447–453.

25. Ross, S. (2021). Machine Learning Is Booming in Medicine. It's Aussi Facing a Credibility Crisis (Stat News). https://www.statnews.com/2021/06/02/machine-learning-ai-methodology-research-flaws/?mod=djemVentureCapitalPro&tpl=vc.

26. Wong, A., Otles, E., Donnelly, J.P., Krumm, A., McCullough, J., DeTroyer-Cooley, O., Pestrue, J., Phillips, M., Konye, J., Penoza, C., and Ghous, M. (2021). External validation of a widely implemented proprietary sepsis prediction model in hospitalized patients. JAMA Intern. Med *181*, 1065–1070.

27. Roberts, M., Driggs, D., Thorpe, M., Gilbey, J., Yeung, M., Ursprung, S., Aviles-Rivero, A.I., Etmann, C., McCague, C., Beer, L., and Weir-McCall, J.R. (2021). Common pitfalls and recommendations for using machine learning to detect and prognosticate for COVID-19 using chest radiographs and CT scans. Nat. Machine Intelligence *3*, 199–217.

28. Kaushal, A., Altman, R., and Langlotz, C. (2020). Geographic distribution of US cohorts used to train deep learning algorithms. JAMA *324*, 1212–1213.

29. Gichoya, J.W., McCoy, L.G., Celi, L.A., and Ghassemi, M. (2021). Equity in essence: a call for operationalising fairness in machine learning for healthcare. BMJ Health Care Inform. *28*, e100289.

30. Sambasivan, N., Arnesen, E., Hutchinson, B., Doshi, T., and Prabhakaran, V. (2021). Re-imagining algorithmic fairness in India and beyond. In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, pp. 315–328.

31. Gaube, S., Suresh, H., Raue, M., Merritt, A., Berkowitz, S.J., Lermer, E., Coughlin, J.F., Guttag, J.V., Colak, E., and Ghassemi, M. (2021). Do as AI say: susceptibility in deployment of clinical decision-aids. NPJ Digital Med. *4*, 1–8.

32. Sambasivan, N., Kapania, S., Highfill, H., Akrong, D., Paritosh, P., and Aroyo, L.M. (2021). "Everyone wants to do the model work, not the data work": data cascades in high-stakes AI. In Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, pp. 1–15.

33. Mitchell, W.G., Dee, E.C., and Celi, L.A. (2021). Generalisability through local validation: overcoming barriers due to data disparity in healthcare. BMC Ophthalmol. *21*, 1–3.

34. Khan, S.M., Liu, X., Nath, S., Korot, E., Faes, L., Wagner, S.K., Keane, P.A., Sebire, N.J., Burton, M.J., and Denniston, A.K. (2020). A global review of publicly available datasets for ophthalmological imaging: barriers to access, usability, and generalisability. Lancet Digital Health *3*, e51–e66.

35. Abebe, R., Aruleba, K., Birhane, A., Kingsley, S., Obaido, G., Remy, S.L., and Sadagopan, S. (2021). Narratives and counternarratives on data sharing in Africa. In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, pp. 329–341.

36. Ismail, A., and Kumar, N. (2018). Engaging solidarity in data collection practices for community health. Proc. ACM Human-Comput. Interact. *2*, 1–24.

37. Bondi, E., Xu, L., Acosta-Navas, D., and Killian, J.A. (2021). Envisioning communities: a participatory approach towards AI for social good. arXiv, 2105.01774.

38. Jacobs, M., He, J., Pradier, M., Lam, B., Ahn, A.C., McCoy, T.H., Perlis, R.H., Doshi-Velez, F., and Gajos, K.Z. (2021). Designing AI for trust and collaboration in time-constrained medical decisions: a sociotechnical lens. In Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, pp. 1–14.

39. Pierre, J., Crooks, R., Currie, M., Paris, B., and Pasquetto, I. (2021). Getting ourselves together: data-centered participatory design research & epistemic burden. In Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, pp. 1–11.

40. Harrington, C., Erete, S., and Piper, A.M. (2019). Deconstructing community-based collaborative design: towards more equitable participatory design engagements. Proc. ACM Human Comput. Interact. *3*, 1–25.

41. Ahuja, A.S. (2019). The impact of artificial intelligence in medicine on the future role of the physician. PeerJ *7*, e7702.

42. Wilson, H.J., and Daugherty, P.R. (2018). Collaborative intelligence: humans and AI are joining forces. Harv. Business Rev. *96*, 114–123.

About the author

Chinasa T. Okolo is a PhD candidate in the department of computer science at Cornell University. Her research interests include AI ethics, computer vision, global health, and information & communication technologies for development. Her research leverages computer vision techniques to improve community health in low-resource regions. Chinasa also studies the applications, implications, and perceptions of AI-enabled technologies deployed throughout the Global South, with a focus on healthcare.