Artificial intelligence (AI) and global health: how can AI contribute to health in resource-poor settings?

Brian Wahl,¹ Aline Cossy-Gantner,² Stefan Germann,² Nina R Schwalbe¹

To cite: Wahl B, Cossy-Gantner A, Germann S, *et al.* Artificial intelligence (AI) and global health: how can AI contribute to health in resource-poor settings? *BMJ Glob Health* 2018;3:e000798. doi:10.1136/bmjgh-2018-000798

Handling editor Seye Abimbola

Received 24 February 2018

Revised 20 July 2018

Accepted 27 July 2018

ABSTRACT

The field of artificial intelligence (AI) has evolved considerably in the last 60 years. While there are now many AI applications that have been deployed in high-income country contexts, use in resource-poor settings remains relatively nascent. With a few notable exceptions, there are limited examples of AI being used in such settings. However, there are signs that this is changing. Several high-profile meetings have been convened in recent years to discuss the development and deployment of AI applications to reduce poverty and deliver a broad range of critical public services. We provide a general overview of AI and how it can be used to improve health outcomes in resource-poor settings. We also describe some of the current ethical debates around patient safety and privacy. Despite current challenges, AI holds tremendous promise for transforming the provision of healthcare services in resource-poor settings. Many health system hurdles in such settings could be overcome with the use of AI and other complementary emerging technologies. Further research and investments in the development of AI tools tailored to resource-poor settings will accelerate realising of the full potential of AI for improving global health.

INTRODUCTION

The field of artificial intelligence (AI) has come a long way since the term was first coined by a group of researchers in 1956.¹ While there are now many AI applications that have been deployed in high-income country contexts, use in resource-poor settings remains relatively nascent. There are signs that this is changing. In 2017, the United Nations (UN) convened a global meeting to discuss the development and deployment of AI applications to reduce poverty and deliver a broad range of critical public services.² More recently, another UN meeting brought together various stakeholders to assess the role AI could play in achieving the Sustainable Development Goals (SDGs).³ⁱ

ⁱ This analysis has been informed by a review of articles published from January 2000 to December 2017 from Medline. Many relevant AI applications for health in resource-poor settings are also described in non-peer reviewed publications. Therefore, we conducted

Summary box

- ▶ Artificial Intelligence (AI) is quickly evolving and is already being used to support and improve health services in many high-income countries.
- ▶ AI holds great promise for improving the delivery of health services in resource-poor settings. Further research and investments are needed to accelerate its deployment in such settings.
- ▶ Deployment and scale up, in particular, in resource-poor settings will also require attention to ethical and legal issues including those related to data privacy.

AI is the branch of computer science that deals with the simulation of intelligent behaviour in computers.⁴ While this definition appears to be straightforward, there is no consensus among experts in the field around what specifically constitutes intelligence. However, some AI experts have proposed that something ‘acts intelligently’ when: (1) what it does is appropriate for its circumstances and its goals; (2) it is flexible to changing environment and changing goals; (3) it learns from experience; and (4) it makes appropriate choices given its perceptual and computational limitations.⁵

TYPES OF AI

There are in general two types of AI. The ability of a machine to represent the human mind and perform any intellectual task that a human can perform is termed artificial general intelligence.⁶ Artificial general intelligence was the initial focus of early AI research and the predominant representation of AI in popular culture.⁷ However, given the challenges and complexity associated with developing this kind of AI, many researchers

supplementary searches using Google and Google Scholar.

© Author(s) (or their employer(s)) 2018. Re-use permitted under CC BY-NC. No commercial re-use. See rights and permissions. Published by BMJ.

¹Spark Street Consulting, New York City, New York, USA

²Fondation Botnar, Basel, Switzerland

Correspondence to

Dr. Brian Wahl;
bpwahl@gmail.com

turned their attention to artificial narrow intelligence: the ability of a machine to perform a single task extremely well. Virtually all contemporary AI health applications are considered artificial narrow intelligence.⁷

No universally accepted classifications of AI subfields relevant to health exist. For that reason, we have chosen to focus on areas that have been repeatedly cited by researchers and are consistent and recurring themes at the Artificial Intelligence in Medicine conference, a prominent biannual meeting of researchers focused on AI applications related to medicine and health.⁸ While additional themes have also been repeatedly discussed in recent years, some of these areas are more upstream (ie, focused on research into foundational AI principles) and thus not included here.⁸ Many of these areas are often used together in a single application.

Expert systems

An expert system, also sometimes referred to as a knowledge-based system, is an AI program that has expert-level competence in solving specific problems. The process of building an expert system is known as knowledge engineering. Expert systems include two primary components: (1) the knowledge base and (2) the reasoning engine. The reasoning engine is often based on a series of complex rules (eg, 'if-then' statements). The information that comprises the knowledge base is almost always incomplete or uncertain. The development of fuzzy logic—a set of mathematical principles for knowledge representation based on probability and uncertainty—has accelerated the evolution of expert systems in recent years. Incorporating fuzzy logic into decision support applications can help better approximate how human would approach complex problems with high degrees of uncertainty. While many fuzzy logic systems are geared towards improving the diagnosis of chronic conditions,^{9–11} researchers in South Africa, for example, used fuzzy logic to improve predictions of cholera outbreaks.¹²

Machine learning

Often conflated with AI, machine learning is a subfield and one application of AI. In practical terms, machine learning is a method for automating data analysis by using algorithms that iteratively identify patterns in data and learn from them. Machine learning applications are generally classified into three broad categories: (1) supervised learning, (2) unsupervised learning and (3) reinforcement learning. Supervised learning uses patterns already identified in data (ie, training data). In contrast, unsupervised learning applications are intended to find and learn from patterns in data. Finally, reinforcement learning is an extension of supervised learning in which 'rewards' and 'punishments' are provided as the application interacts with a dynamic environment. Data mining is related to unsupervised machine learning and involves identifying patterns in large datasets. The differences are subtle but important.

Natural language processing (NLP)

NLP is another subfield of AI that aims to bridge the divide between the languages that humans and computers use to operate. By using algorithms that allow machines to identify key words and phrases in natural language corpora (ie, unstructured written text), AI applications are able to determine the meaning of text. The field of NLP has quickly evolved to focus on managing and processing information from large datasets. Topic modelling is an approach to NLP that seeks to automatically identify the topics covered in documents by inferring relationships among prominently featured words.

Automated planning and scheduling

Automated planning and scheduling is a relatively nascent branch of AI focused on organising and prioritising the activities required to achieve a desired goal. It is also sometimes referred to as AI planning. In many such applications, optimised processes are executed by other AI applications or autonomous robot; this is not always the case. Some automated planning applications can be used to improve the efficiency of human procedures.

Image and signal processing

AI can also be used to process large amounts of data from images and signals (ie, information about the attributes of a particular physical phenomenon). Data produced by motion and sound are common examples of signals. Steps in image and signal processing algorithms typically include signal feature analysis and data classification using tools such as artificial neural networks (ANNs). ANNs are computing systems based on the networks that comprise animal brains. In recent years, the availability of physiological instruments, such as wearable fitness trackers, that provide biomedical measurements and the processing power from mobile phones have enabled virtually real-time signal data collection that can be processed using AI applications.

COMPLEMENTARY ADVANCES IN DIGITAL HEALTH TECHNOLOGIES IN RESOURCE-POOR SETTINGS

High-income countries are already benefiting from integrating AI into their healthcare ecosystems. One analysis has found that the use of AI applications could result in approximately \$150 billion in saved healthcare costs annually by 2026 in the USA.¹³ There are many reasons to be optimistic that AI could also prove transformative for public health in resource-poor countries. The large volume of data being generated has created a plethora of opportunities for applying AI to improve individual and population health. This is particularly true in resource-poor settings where there has been strong mobile phone penetration, developments in cloud computing and substantial investments in digitising health information and introducing mobile health (mHealth) applications. As a result, many places now have the necessary basics in place to initiate meaningful applications of AI. Advances

in AI could also help expand and strengthen the impact of these and other digital health technologies in such settings.

Health informatics and electronic medical records (EMRs)

Health informatics describes the acquisition, storage, retrieval and use of healthcare information to improve patient care across interactions with the health system. Health informatics can help shape public health programmes by ensuring that critical information is available for making sound policies and programme decisions. EMRs, which are digital versions of patient and population health information, are an important source of data for health informatics. Their use has become much more prevalent in low-resource settings, which in an era of networked computers, has expanded potential applications of AI to improve public health informatics and decision making. OpenMRS is one example of an EMR platform that is currently being used in more than 15 African countries. OpenMRS was used in Kenya to build and implement an EMR system called Bora to improve maternal and child health and improve HIV treatment in rural areas. Researchers found that this system helped to increase the completeness of data collected and close critical gaps in care.^{14 15} DHIS2 is another open source EMR platform used for collecting, validating, analysing and presenting aggregate and patient-based statistical data. It is now used in more than 40 countries of Asia, Africa and Latin America.¹⁶

Cloud computing

The expansion of cloud computing has led to the expansion of AI applications for health. Cloud computing refers to the use of a network of remote servers to store, manage, access and process data rather than a single personal computer or hard drive. Businesses have been quick to adopt cloud computing because of the many advantages it offers over maintaining inhouse information technology (IT) systems, including increased reliability and substantial cost savings. Through companies that provide cloud computing services, resource-strapped organisations, particularly those in low- and middle-income countries (LMICs), are now able to access computing power that would have been unattainable previously. While EMRs can be maintained in the cloud with adequate privacy and security precautions, cloud computing can be used with a multitude of data related to public health. One appealing advantage of cloud computing is that it enables the implementation of applications that rely on IT infrastructure in settings where it does not currently exist. For example, researchers recently tested a cloud computing application using patient data that aimed to improve interactive voice response telephone calls for managing non-communicable diseases in Honduras. The researchers found that the intervention was effective despite the lack of adequate IT infrastructure in the country.¹⁷

Mobile health

mHealth uses mobile and wireless technologies to achieve health objectives. The rapid availability and expansion of mobile phones in low-income countries has created several opportunities for using these technologies to support health efforts. Mobile phones have been used by community health workers (CHWs) to improve the provision of health services within resource-poor settings. In Tanzania, a non-governmental organization has developed a mobile phone-based tool that assists CHWs deliver targeted information to pregnant women and caregivers of young children.¹⁸ The overall impact of the tool is currently being assessed. Mobile phones have also been used to communicate health information to patients in resource-poor settings when face-to-face interactions are not feasible. The use of short message service to address demand-side barriers to vaccination and improve immunisation coverage has been thoroughly documented through randomised controlled trials in settings such as Kenya¹⁹ and the USA.²⁰

AI APPLICATIONS IN RESOURCE-POOR SETTINGS

Little has been documented in the academic literature on AI applications for health in resource-poor settings. However, this should not be taken as an indicator of the current activities in this area. There are many examples of promising technologies covered in the media and at conferences. The 'AI for Development' community has published several observations of how AI has been implemented in resource-poor settings. A major lesson from the experience of those working on AI in resource-poor settings is that 'AI should build intelligence into existing systems and institutions rather than starting from scratch or hoping to replace existing systems, however broken'.²¹ The success of AI applications requires knowledge of local markets, clear usability requirements and access to adequate training data via field testing.²¹ Examples of how AI is currently being deployed are noted below. While many of these interventions have not yet been robustly evaluated, they do provide insight into how AI subfields are being applied and their potential beyond current pilots.

In resource-poor settings, expert systems can be used to support health programmes in several ways. First, medical expert systems can support physicians in diagnosing patients and choosing treatment plans as is done in high-income countries. For some conditions, they can act in place of a human expert if one is not readily available, which is often the case in poor communities. In this way, expert systems have already been deployed in resource-poor settings. Birth asphyxia is not always predictable or preventable. Cases often require immediate, skilled resuscitation in the delivery room. Researchers from Brazil and USA developed a fuzzy logic-based expert system to predict birth asphyxia for use in developing countries. The system relates maternal medical, obstetric and neonatal characteristics to the clinical conditions

of the newborn to predict the need for resuscitation. Researchers found that the application was 77% sensitive and 95% specific for identifying the need for birth asphyxia at tertiary perinatal medical centres.²²

AI is already being used to predict, model and slow the spread of disease in epidemic situations around the world, including in resource-poor settings. Dengue fever is a vector-borne disease that has spread rapidly around the globe in recent years. About half of the world's population is currently at risk.²³ Researchers have developed a machine learning tool to identify weather and land-use patterns associated with dengue fever transmission in Manila. The machine learning algorithm has learnt over many iterations how to fine-tune its model to predict dengue occurrence with increasing accuracy.²⁴ Researchers are now working with the Philippine government to expand the programme.

There are currently two ways in which NLP is explicitly being used to address health challenges in resource-poor settings. First, it is being used for surveillance and outbreak predictions using data from electronic health records and online media and social media sources. Global Health Monitor is one example of an online system for detecting and mapping infectious disease outbreaks. The tool works by identifying English-language news stories and classifying them for topical relevance. Relevant articles are then plotted on a map using geo-coding information, which can help epidemiologists and programme managers monitor the spread of diseases.²⁵

NLP is also being used to analyse unstructured text in the medical literature and EMRs to support clinical decision making or track population health behaviour. Using data from the web, for example, NLP has been applied to a wide range of public health challenges, from improving treatment protocols to tracking health disparities.^{26 27} NLP and machine learning are also being used to guide cancer treatments in low-resource settings including in Thailand, China and India.²⁸ Researchers trained an AI application to provide appropriate cancer treatment recommendations by giving it descriptions of patients and telling the application the best treatment options. The AI application uses NLP to mine the medical literature and patient records—including doctor notes and lab results—to provide treatment advice. When examining different patients, this application agreed with experts in more than 90% of patients in one study and 50% in another.²⁸

AI planning is also already being applied to improving the provision of primary healthcare services in resource-poor settings. These tools—versions of which are under development or are being piloted—are being applied to improve the efficiency of programmes ranging from immunisation, to supply chain, or referral services within a multitiered system. CHWs often visit several households per day—often separated by long distances—where they provide health information, dispense some medicines or perform limited medical procedures. A partnership comprising researchers and a social enterprise have been

developing an AI planning application for optimising CHW scheduling in communities in Africa.²⁹ As AI planning, like many AI disciplines, relies on the availability of relevant high-quality data, advancements in mobile connectivity and digitised health system performance and resource allocation data present additional opportunities for AI planning in resource-poor settings.

Signal processing is another related area that could be buoyed by the rapid expansion of mobile devices that can capture and transmit signals and the emergence of cloud computing. Specific opportunities in resource-poor settings currently focus on signals that can be collected with mobile phones or single-function, add-on devices that link with mobile phones or multipurpose instruments that can collect multiple kinds of digital data. Such signal processing opportunities in resource-poor settings are promising, particularly when paired with machine learning principles and cloud computing. For example, signal processing and machine learning have been used in Nigeria to predict birth asphyxia using a mobile phone.³⁰ The application uses the birth cry of children collected using a mobile device to identify children likely to be experiencing birth asphyxia in resource-poor settings. Applications that collect and transmit signals collected by sensors that continuously record whatever is being monitored (eg, activity trackers) are further upstream in such settings.

As resource-poor settings become more connected and the data they produce are of higher quality, the ability of AI to address health challenges will likely expand. Furthermore, as countries develop, several of the emerging digital technologies described here could be employed to monitor patient populations and medical data to recognise patterns that could potentially identify a pending outbreak or other public health emergencies. This would allow public health agencies, such as the WHO, to investigate and monitor in real-time to model cause-and-effect relationships that could mitigate the progression of an epidemic.

UNANSWERED QUESTIONS AND REMAINING CHALLENGES

Sheikhtaheri *et al*³¹ describe many challenges with designing and implementing expert systems for supporting clinical decision making. In general, they conclude that the successful implementation of any expert system requires a clear definition of the clinical problem to be addressed. Building and updating the knowledge base is challenging in the best of circumstances and would be compounded in resource-poor settings. Many expert systems also lack an accuracy tracking mechanism, which could undermine the trust of clinicians and patients. Sheikhtaheri *et al*³¹ also raise interesting questions of whether all disease domains require expert systems and how various AI systems should be integrated. They suggest that the lack of answers for these questions has prevented expert systems from advancing from research to implementation.

Supervised machine learning applications require high-quality datasets that can be used to train machine learning algorithms to identify risk factors or make disease diagnoses.³¹ For many diseases and conditions relevant to resource-poor settings, such datasets can be difficult and time-consuming to collect.³¹ In addition, better diagnosis does not equate to access to appropriate or quality treatment options. While remote diagnostics and machine learning applications might help to identify diseases, there is no guarantee that the condition identified can be treated in any given setting. As such, the principle of 'do no harm' and the ethics around providing treatment following testing for and confirmation of disease ('test and treat') when deploying AI applications are relevant as much as in any other context.

Carrel *et al*³² describe some of the challenges associated with adapting clinical NLP systems to diverse healthcare settings. Using colonoscopy screening as an example, they highlight the substantial resources necessary to compile the natural language corpora, address different record structures and deal with idiosyncratic linguistic content. These challenges would likely be multiplied when considering healthcare settings in low-income countries. Countries in low-resource settings sometimes maintain hand-written health records in local languages. Therefore, building the natural language corpora could require substantial effort. The WHO has advocated for the adoption of standardised medical terminologies or the development of local data dictionaries to address some of these challenges.³³

Substantial data are necessary to build and implement automated planning and scheduling applications.²⁹ Compiling such data in resource-poor settings is often difficult and time-consuming. Brunskill and Lesh describe the extensive data collection involved in laying the groundwork for the development of an improved CHW schedule in sub-Saharan Africa.²⁹ Furthermore, the effectiveness of automated planning and scheduling applications will depend largely on the quality of data used to develop the application. High-quality health systems data are currently difficult to collect in many resource-poor settings. Such challenges could slow the development of automated planning and scheduling in the settings where they would be most needed.

While internet connectivity is improving throughout the world, some resource-poor settings remain without access to the substantial bandwidth necessary to upload very large signal datasets to the cloud. Some applications, including mHealth tools used by CHW, have the ability to work offline and sync with remote databases when the bandwidth is sufficient. Storing such data locally could also require substantial investment in IT infrastructure. While device prices are currently also a barrier, these are predicted to come down over time as companies take advantage of economies of scale.

There are also environmental challenges that need to be considered. The use of AI in resource-poor

settings require a strong understanding of local social contexts, infrastructure requirements and additional related infrastructure needs, including IT, communications networks and platforms for delivering primary health services. Many AI applications depend on the availability of strong electronic health record systems, which require substantial investment to put into place. In addition, AI applications will have limited impact if they do not effectively integrate languages and scripts used in the electronic health records of many developing countries.

In high-income countries, discussions around the ethics of electronic health records and AI have focused largely on privacy, confidentiality, data security, informed consent and data ownership. Most of these same considerations apply to resource-poor settings. However, the relevance of these issues varies depending on differences in culture, literacy, patient-provider relationships, available IT infrastructure and regulatory issues in LMICs. One proposed approach for maintaining secure and transparent health records is the use of 'blockchain', the distributed ledger system known primarily for its use by crypto-currencies.

In addition, some experts have raised concerns that some AI applications can potentially exacerbate inequities, including those related to ethnic, socioeconomic and gender. They note that cultural prejudices can be reflected in data, algorithms and other aspects of AI design.³⁴ One recent report found that an application that uses arrest records, postal codes and socioeconomic data to assess the risk of recidivism in US courts was biased against black citizens.³⁵ These challenges are compounded by the fact that many AI algorithms are a 'black box' and are therefore less likely to be assessed for bias. However, some researchers are working to assess biases by testing how well they predict by randomly changing key variables for individuals for whom AI applications are attempting to make a prediction.³⁶ For AI to benefit all, including those in resource-poor settings, these biases need to be considered in the design of such applications. In health, ensuring that more women and those resource-poor communities are involved in the development of AI applications will also help to reduce such biases.

The generation of large amounts of data naturally raises questions around the ownership of the data and who can access which specific data for research or commercial purposes. While there have been some initial efforts to address this, including the *Data Sharing Principles in Developing Countries*, which was put forward in Nairobi in 2014, there has yet to be widespread adoption.³⁷ The underlying concept that initiated the development of these principles is that data generated with public funds should be viewed as a public good. Establishing a repository for large amounts of data for global health would help in ensuring that such data are made available as a global public good. Many LMICs have signed agreements that data generated with use of public funds should be freely

available.³⁷ However, both from patient and developer perspectives, privacy laws and data access and ownership agreements are perceived to be potential threats to successful AI applications and they should therefore be monitored closely by groups working to develop such applications for particular contexts. Applications using new technologies like blockchain, may also help resolve some of the current concerns.

CONCLUSION

AI holds tremendous promise for transforming the provision of healthcare services in resource-poor settings. Many of the health systems hurdles in such environments could be addressed and overcome using AI supported by other technological developments and emerging fields. The ubiquitous use of smartphones, combined with growing investments in supporting technologies (eg, mHealth, EMR and cloud computing), provide ample opportunities to use AI applications to improve public health outcomes in low-income country settings. While we have provided several examples of how AI is already being applied with the aim to improve health outcomes in low-income countries, there are certainly many other AI applications already being implemented and surely there will be more in coming years.

Moving from pilot to scale in these settings will require addressing several issues and must be grounded in the experience of the beneficiaries of these powerful tools. That means using human-centred design when developing and implementing new AI applications. It also means considering legal and ethical questions through a human rights lens that includes privacy, confidentiality, data security, ownership and informed consent. Effective implementation will also require understanding the local social, epidemiological, health system and political contexts. Furthermore, wide-scale deployment will need to be guided by a robust research agenda. Although not a panacea, AI is one of several tools that could help in achieving the health-related targets set out in the SDGs, particularly those related to providing universal health coverage.

Contributors BW and NRS designed and conducted the analysis. They also prepared the first draft of the manuscript. SG and AC-G conceptualised the project and provided substantial input to the analysis and edits to the manuscript.

Funding Fondation Botnar.

Competing interests None declared.

Patient consent Not required.

Provenance and peer review Not commissioned; externally peer reviewed.

Data sharing statement No additional data are available.

Open access This is an open access article distributed in accordance with the Creative Commons Attribution Non Commercial (CC BY-NC 4.0) license, which permits others to distribute, remix, adapt, build upon this work non-commercially, and license their derivative works on different terms, provided the original work is properly cited, appropriate credit is given, any changes made indicated, and the use is non-commercial. See: <http://creativecommons.org/licenses/by-nc/4.0/>.

REFERENCES

- Knapp S. Artificial intelligence: past, present, and future. 2006. <http://www.dartmouth.edu/~vox/0607/0724/ai50.html>
- AI for Good Global Summit, Geneva. 2017. <http://www.itu.int/en/ITU-T/AI/Pages/201706-default.aspx>
- Looking to future UN to consider how artificial intelligence could help achieve economic growth and reduce inequalities. <http://www.un.org/sustainabledevelopment/blog/2017/10/looking-to-future-un-to-consider-how-artificial-intelligence-could-help-achieve-economic-growth-and-reduce-inequalities/2017>
- Malley CS, Kuylensstierna JC, Vallack HW, et al. Preterm birth associated with maternal fine particulate matter exposure: a global, regional and national assessment. *Environ Int* 2017;101:173–82.
- Poole DL, Mackworth AK. *Artificial Intelligence: foundations of computational agents*: Cambridge University Press, 2010.
- Goertzel B, Pennachin C. *Artificial general intelligence*: Springer, 2007.
- Pennachin C, Goertzel B. Contemporary approaches to artificial general intelligence. *Artificial general intelligence: Springer* 2007:1–30.
- Peek N, Combi C, Marin R, et al. Thirty years of artificial intelligence in medicine (AIME) conferences: a review of research themes. *Artif Intell Med* 2015;65:61–73.
- Fathi-Torbaghan M, Meyer D. MEDUSA: a fuzzy expert system for medical diagnosis of acute abdominal pain. *Methods Inf Med* 1994;33:522–9.
- Lee CS, Wang MH. A fuzzy expert system for diabetes decision support application. *IEEE Trans Syst Man Cybern B Cybern* 2011;41:139–53.
- Abbod MF, Catto JW, Linkens DA, et al. Application of artificial intelligence to the management of urological cancer. *J Urol* 2007;178(4 Pt 1):1150–6.
- Fleming G, Merwe Mvander, McFerren G. Fuzzy expert systems and GIS for cholera health risk prediction in southern Africa. *Environmental Modelling & Software* 2007;22:442–8.
- Collier M, Fu R, Yin L. *Artificial intelligence: healthcare's new nervous system*: Accenture, 2017.
- Haskew J, Rø G, Saito K, et al. Implementation of a cloud-based electronic medical record for maternal and child health in rural Kenya. *Int J Med Inform* 2015;84:349–54.
- Haskew J, Ro G, Turner K, et al. Implementation of a Cloud-Based Electronic Medical Record to Reduce Gaps in the HIV Treatment Continuum in Rural Kenya. *PLoS One* 2015;10:e0135361.
- DHIS. <https://www.dhis2.org/overview>.
- Piette JD, Mendoza-Avelares MO, Ganser M, et al. A preliminary study of a cloud-computing model for chronic illness self-care support in an underdeveloped country. *Am J Prev Med* 2011;40:629–32.
- World Vision, International. mHealth: Tanzania. <https://www.wvi.org/health/mhealth-tanzania>
- Gibson DG, Ochieng B, Kagucia EW, et al. Mobile phone-delivered reminders and incentives to improve childhood immunisation coverage and timeliness in Kenya (M-SIMU): a cluster randomised controlled trial. *Lancet Glob Health* 2017;5:e428–38.
- Stockwell MS, Kharbanda EO, Martinez RA, et al. Effect of a text messaging intervention on influenza vaccination in an urban, low-income pediatric and adolescent population: a randomized controlled trial. *JAMA* 2012;307:1702–8.
- Remembering the Past for Meaningful AI-D*. AAAI Spring Symposium: Artificial Intelligence for Development, 2010.
- Reis MA, Ortega NR, Silveira PS. Fuzzy expert system in the prediction of neonatal resuscitation. *Braz J Med Biol Res* 2004;37:755–64.
- Dengue and severe dengue. 2017. <http://www.who.int/mediacentre/factsheets/fs117/en/2017>
- Hornyak T. Mapping Dengue Fever Hazard with Machine Learning. *Eos* 2017. <https://eos.org/articles/mapping-dengue-fever-hazard-with-machine-learning> (updated 14 June 2017).
- N.: Global Health Monitor-a web-based system for detecting and mapping infectious diseases. In: *Proceedings of the International Joint Conference on Natural Language Processing*. Citeseer, 2008:951–6.
- International Workshop on Knowledge Representation for Health Care. *A public health surveillance platform exploiting free-text sources via natural language processing and linked data: application in adverse drug reaction signal detection using pubmed and twitter*: Springer, 2016.
- Wieland ML, Wu ST, Kaggal VC, et al. Tracking health disparities through natural-language processing. *Am J Public Health* 2013;103:448–9.

28. Swetlitz I. Watson goes to Asia as hospitals use supercomputer for cancer treatment: PBS NewsHour. 2016. <http://www.pbs.org/newshour/rundown/hospitals-use-supercomputer-cancer-treatment/> 2017
29. Routing for Rural Health. *Optimizing Community Health Worker Visit Schedules*. AAAI Spring Symposium: Artificial Intelligence for Development, 2010.
30. Onu CC, Udeogu I, Ndiomu E, *et al*. Ubenwa: Cry-based Diagnosis of Birth Asphyxia. *arXiv preprint arXiv* 2017;171106405.
31. Sheikhtaheri A, Sadoughi F, Hashemi Dehaghi Z. Developing and using expert systems and neural networks in medicine: a review on benefits and challenges. *J Med Syst* 2014;38:110.
32. Carrell DS, Schoen RE, Leffler DA, *et al*. Challenges in adapting existing clinical natural language processing systems to multiple, diverse health care settings. *J Am Med Inform Assoc* 2017;24:986–91.
33. Organization WH. *Electronic health records: manual for developing countries*. Manila: WHO Regional Office for the Western Pacific, 2006.
34. Caliskan A, Bryson JJ, Narayanan A. Semantics derived automatically from language corpora contain human-like biases. *Science* 2017;356:183–6.
35. Angwin JLK, Mattu S, Kirchner L. *Machine bias*: ProPublica, 2016.
36. IEEE Symposium. *Algorithmic transparency via quantitative input influence: Theory and experiments with learning systems. Security and Privacy (SP)*: IEEE, 2016.
37. *Data sharing principles in developing countries: the Nairobi data sharing principles*. Nairobi, 2014.