Research Paper

Claire Muñoz Parry and Urvashi Aneja
Centre for Universal Health & Asia-Pacific Programme | July 2020

Artificial Intelligence for Healthcare
Insights from India

CHATHAM HOUSE

Contents

Summary

- AI, the use of coded computer software routines with specific instructions to perform tasks for which a human brain is considered necessary, is providing the healthcare sector with new advances that are being hailed as game changers.
- The risk and challenges of integrating AI into healthcare are closely related to the use of the data needed to feed AI systems. Issues around quality, safety, governance, privacy, consent and ownership must all be properly addressed. A lack of explainability, as it is almost impossible to understand how AI arrived at a specific decision, also points to a potential lack of trust in AI systems.
- India provides a case study of how a country is actively promoting the use of AI to address healthcare needs. However, the deployment of AI in India is still at a very nascent stage, particularly for clinical interventions.
- The challenge of delivering quality healthcare at scale presents a strong case for developing AI-based solutions for healthcare in India. However, a complex health landscape involving numerous stakeholders, competing priorities, entrenched incentive systems and institutional cultures give rise to a range of challenges and risks across the stages of development, adoption and deployment.
- The quality of digital infrastructure, affordability, and variable capacity among states and medical professionals are together likely to result in adoption of AI applications primarily by India's well-established private hospitals. This in turn could result in new inequities in access to quality healthcare.
- The effectiveness of AI systems will depend on accurate problem identification and solution matching. Currently, there is a risk that solutions are being technology-led rather than problem-led, and as a result are often blind to particular contextual needs or constraints.

1. Introduction

Artificial intelligence (AI), the technology that has captivated multiple sectors, is being hailed as a tool that will help provide access for all to quality medical care, including through the development and improvement of diagnostics, personalized medical care, the prevention of illnesses and the discovery of new treatments. Within the next five years, the use of AI in medicine is expected to increase tenfold (Perry, 2016).

AI can be defined as the use of coded computer software routines (algorithms) with specific instructions to perform tasks for which a human brain is normally considered necessary. Such software can help people understand and process language, recognize sounds, identify objects and make use of learning patterns to solve problems. Machine learning (ML) is a way of continuously refining an algorithm. The refinement process involves the use of large amounts of data and is done automatically, allowing the algorithm to change with the aim of improving the precision of the artificial intelligence (Zandi, 2019). Put simply, AI enables computers to model intelligent behaviour with minimal human intervention, and has been shown to outperform human beings at specific tasks. In 2017, for instance, it was reported that deep neural networks (a branch of AI) had been used successfully to analyse skin cancer images with greater accuracy than a dermatologist, and to diagnose diabetic retinopathy (DR) from retinal images (*The Lancet*, 2017).

However, the definition of AI is evolving. As well as the more technical definition given above, AI is also perceived as something resembling human intelligence, aspiring to exceed the capabilities of any of the individual technologies. It is conceived as a technology interaction that gives a machine the ability to fulfil a function that 'feels' human. The ability of a machine to perform any task that can be achieved by a human has been termed Artificial General Intelligence (AGI). AGI systems are designed with the human brain as a reference. However, AGI has not yet been achieved; experts recently forecast its emergence by 2060 (Joshi, 2019).

The case for examining the potential opportunities and risks of implementing AI systems for healthcare purposes has been given new importance by the outbreak of the COVID-19 virus, which plunged the world into a public health crisis of unprecedented proportions from early 2020. AI systems could help overburdened health administrations to plan and rationalize resources, and to predict new COVID-19 hotspots and transmission trends, as well as provide a critical tool in the search for drug treatments or vaccines. However, as governments around the world scramble to adopt technological solutions (many of which rely on ML systems) to help them contain and mitigate the crisis, questions around the ethics and governance of AI are arising with equal urgency. There are already growing concerns about how the current crisis is going to expand governments' surveillance capacities, as well as accentuating the power and influence of so-called 'big tech' companies. These concerns are particularly acute for developing countries such as India. In such countries, weak public health infrastructure is increasing the appeal of AI-based solutions, even while the normative and regulatory frameworks required to steer AI trajectories are weak and underdeveloped.

This paper describes some of the main opportunities and challenges of using AI in healthcare. It then turns to a case study of the use of AI for healthcare purposes in India, discussing key applications, challenges and risks in this context.

2. AI and Healthcare

What is AI?

Artificial intelligence for health includes ML, natural language processing (NLP), speech recognition (text-to-speech and speech-to-text), image recognition and machine vision, expert systems (a computer system that emulates the decision-making ability of a human expert), robotics, and systems for planning, scheduling and optimization.

ML is a core component of AI that provides systems with the ability to automatically learn and improve without being explicitly programmed. In fact, there cannot be AI without ML. Computer programmes access data and use it with the aim of learning without human intervention or assistance, and adjust actions accordingly (Expert System, 2017). Deep learning (DL), a type of ML, is inspired by the human brain, and uses multi-layered neural networks to find complex patterns and relationships in large datasets that traditional ML may miss (Health Nucleus, undated).

NLP is a subfield of AI that helps computers understand, interpret and manipulate human language. It draws from many disciplines, including computer science, linguistics, information engineering and computational linguistics, in pursuit of filling the gap between human communication and computer understanding (SAS, undated).

Speech recognition is the ability of a machine or programme (a mix of software and hardware) to identify words and phrases in spoken language and convert them to a machine-readable format, and vice versa. It is also known as automatic voice recognition (AVR) or voice-to-text.

The promise of AI for healthcare

The World Economic Forum has proposed four ways in which AI can make healthcare more efficient and affordable: enabling tailored treatment plans that will improve patient outcomes, and therefore reduce the cost associated with complications arising from treatment; permitting better and earlier diagnosis that reduces human error; enabling accelerated drug development; and empowering patients to take a more active role in managing their health (World Economic Forum, 2018). One of AI's main attractions is the potential savings it could bring to the healthcare sector. According to a study by Accenture, when combined, key clinical AI applications could create $150 billion in annual savings for the US healthcare economy by 2026. AI can help minimize preventable and rectifiable system inefficiencies (such as over-treatment, improper care delivery or, indeed, care delivery failures), ensuring substantially more streamlined and cost-effective health ecosystems (Accenture, 2017).

A further benefit of the application of AI to healthcare settings would be the liberation of health workers from hours of mundane data work. They would thus be able to focus more on patient care, leaving to technology the task of examining and analysing clinical data. This, for example, would allow healthcare practitioners to assess patients with greater precision, which in turn would translate into faster and more accurate diagnoses. AI can provide a diagnosis that would have taken a doctor (or a team of doctors) many hours to reach. It can process a huge volume of medical images and scans in a fraction of the time that the same task would take a human expert. In this respect, AI is already revolutionizing the field of radiology by improving workflows, diagnostic and imaging assistance.

Likewise, the use of AI for administrative purposes will free up resources that can be focused on delivery of care, the creation of new drugs and therapies, and the conducting of research to eradicate diseases. Doctors, nurses and other healthcare workers will be relieved of laborious tasks that contribute to burnout, thereby also reducing human errors in the practice of medicine (Ash, Petro and Rab, 2019). NLP, for example, is used to analyse unstructured clinical notes, prepare reports and transcribe interactions with patients. Robotic process automation (which in fact involves computer programmes hosted on servers, rather than robots) is used for repetitive tasks such as prior authorization (required by some health insurance schemes), updating patient records or billing (Davenport and Kalakota, 2019).

> Robotics outcomes include a 21 per cent reduction in length of stay. The value of robotic solutions will increase further as their development and use progresses to a greater diversity of surgeries.

At least for high-income countries, one of several AI applications that has garnered significant interest is robot-assisted surgery. ‘Cognitive robotics’ can integrate information from pre-operative medical records with real-time operating metrics to physically guide and enhance the physician’s instrument precision. The technology incorporates data from actual surgical experiences to inform new, improved techniques and insights. Robotics outcomes include a 21 per cent reduction in length of stay (Accenture, 2017). The value of robotic solutions will increase further as their development and use progresses to a greater diversity of surgeries.

According to Accenture, the benefits from AI accrue incrementally, from automated operations, precision surgery and preventive intervention (thanks to predictive diagnostics), and within a decade they are expected to fundamentally reshape the healthcare landscape (*Forbes*, 2019).

Low- and middle-income countries (LMICs) will, it is hoped, in time have access to costly and highly sophisticated AI applications such as robot-assisted surgery. Currently, healthcare systems in low-resource settings are dealing with shortages of workers, medical equipment and other infrastructure. But AI tools could optimize existing resources and help overcome workforce resource shortages, while also greatly improving healthcare delivery and outcomes in ways never previously imagined (USAID, 2019). The greatest near-term value of AI in LMICs is considered by some to be in squeezing more value out of available data through ML. Some key applications of AI for health in LMICs are expected to increase access to healthcare as well as enhancing its quality. Such programmes focus on monitoring and assessing population health, and targeting public health interventions to better effect; enabling frontline health workers – including community health workers – to better serve their patients, using AI-powered tools such as mobile phone apps; developing virtual ‘health assistants’ that are able to coach patients in managing their conditions or to advise them when to seek care; and developing tools to help doctors diagnose and treat their patients (ibid.).

Through the use of data science (a multidisciplinary field that uses scientific methods, processes, algorithms and systems to extract knowledge and insights from structured and unstructured data), AI is being used in a diverse set of therapy areas, including wellness and lifestyle management, diagnostics, wearables and virtual assistants; it is also being used for disease surveillance to predict, model and slow the spread of disease in epidemic situations, including in resource-poor settings (ibid.).

A recent example is an ML tool that has been used in the Philippines to identify weather and land-use patterns associated with the transmission of dengue fever, a mosquito-borne disease that has spread rapidly around the globe in recent years (Wahl et al., 2018). The AI machine, produced by AIME (Artificial Intelligence in Medical Epidemiology, a US-based company), is able to predict dengue occurrence with increasing accuracy. AIME's technology has been deployed in Rio de Janeiro, Singapore, the Dominican Republic and two states in Malaysia. The platform provides users with three months' advance notice of the exact geolocation and date of the next dengue outbreak. Its customized analytics platform also makes sense of its users' public health data and provides time charts, historic mapping of diseases and 'rumour reports' from social media (World Wide Web Foundation, 2017).

At a policy level, these new sciences offer the possibility of supporting health policy decision-making, a better integration of healthcare with other sectors, and substantial time and efficiency savings in undertaking research and driving quality improvement initiatives.

At a policy level, these new sciences offer the possibility of supporting health policy decision-making, a better integration of healthcare with other sectors, and substantial time and efficiency savings in undertaking research and driving quality improvement initiatives (Colclough et al., 2018). At a time when transformation in health systems is increasingly needed to deal with the new challenges of a growing, ageing population that suffers from a number of medical conditions, the use of AI to process the datasets associated with these cases promises to be invaluable.

It is precisely AI's ability to carry out speedy processing and analysis of datasets that is one of its key strengths. With more countries perfecting the use of health informatics and electronic medical records (EMR), AI will become increasingly useful. In India, 30–60 per cent of the population have declared that they would want their health data to be shared to improve care delivery, to permit research to be conducted and to inform health planning (ibid.). In Kenya, an open EMR platform has contributed to improving child and maternal health and HIV/AIDS treatments in rural areas by helping to achieve a more complete data collection. The cloud-based EMR system was used in western Kenya in 2013. Results of a study showed that implementation of the system resulted in a 42.9 per cent improvement in the completeness of data (including screening for hypertension, tuberculosis, malaria, HIV/AIDS status, or antiretroviral therapy (ART) status of HIV-positive women) (Haskew et al., 2015).

The use of NLP technologies allows machines to identify key words and phrases, and enables them to determine the meaning of text. NLP algorithms are used, for example, to simplify clinical documentation and enable voice-to-text dictation. These technologies are in increasing demand by healthcare providers who are challenged by electronic health record (EHR)[1] overload, as they allow them to interact with patients and produce accurate records of consultations without having to type at the same time. Both Google and Amazon are exploring, respectively, how to turn Google Home and Alexa, their popular ambient home computing devices, into innovative healthcare 'helpers'.

[1] There is a difference between electronic medical records (EMR) and electronic health records (EHR). EMR contain the medical and treatment history of patients, while EHR focus on the total health of the patient – going beyond standard clinical data collected in the provider's office and inclusive of a broader view on a patient's care (Garrett and Seidman, 2011).

In May 2018, for instance, there were reports that Amazon was planning to leverage Alexa for chronic disease management and home care (Health IT Analytics, undated). NLP is also being used to guide cancer treatments in low-resource settings, including in Thailand, China and India, where AI mines the medical literature and patient records – including doctors' notes and lab results – to provide treatment advice (Wahl et al., 2018).

As health-focused IT tools such as NLP become more advanced, the potential of using them to improve the care continuum can only become greater. In resource-poor environments, AI and its complementary technologies could help to overcome hurdles in healthcare systems. High levels of mobile phone penetration, developments in cloud computing, substantial investments in digitizing health information and the introduction of mobile health (mHealth) applications are providing a wide range of opportunities for AI to improve individual and population health.

Ethical, legal and other challenges in the use of AI in healthcare

AI applications, along with technologies such as big data and robotics, are expected to have transformational and disruptive potential within the healthcare sector – across various areas such as hospitals and hospital management, pharmaceuticals, mental health and well-being, insurance, and predictive and preventive medicine. However, these applications introduce new risks and challenges that will require policy and institutional frameworks to guide AI design and use. This paper focuses mostly on challenges at the individual level.

With the increasing availability of health-related data, and the use of AI to analyse such data for medical purposes, ethical, technical and resource-related questions will need to be answered. There are quality, safety, governance, privacy, consent and ownership challenges that are still under-addressed. There is also concern among those examining AI design and use that there is a need for humans to understand why and how AI arrived at a specific decision. The processes AI follows, and the speed with which it deals with large amounts of information, are beyond human perception. Many of the algorithms created by ML cannot be easily examined, and it is impossible to understand specifically how and why AI arrived at a specific conclusion. A lack of explainability and trust in AI processes hampers the ability to fully trust AI systems (Schmelzer, 2019).

In LMICs, some of the challenges of integrating AI into healthcare systems relate to the hurdles of scaling digital health technologies. Other challenges are linked to the fact that '[…] LMIC governments lack the resources and technological capabilities to create consistent policies on population health, such as disease burden analysis and monitoring and treatment protocols for use, across their various regions or states. This creates a barrier for AI tools for population health to scale at a national level.' (USAID, 2019) In terms of quality, AI requires high-quality data in order to produce coherent results. In low-resource settings, this is not always available. A strong digital health infrastructure is required to operate AI tools. Low EMR adoption rates (estimated at less than 40 per cent in LMICs) constitute one of the barriers to feeding AI machines with the necessary historic and real-time patient data (World Bank, 2019; USAID, 2019). Even in high-income countries, the quality of data is a factor determining the speed at which AI tools are put into use. The average UK hospital, for instance, has hundreds of different systems that are not integrated with each other. There is a need for 'an interconnected data infrastructure with fast, reliable and secure interfaces, international standards for data exchange as well as medical terminologies that define unambiguous vocabularies for the communication of medical information' (Lehne et al., 2019).

Protection of citizens' health data is a key area of responsibility for those handling sensitive data for AI purposes. Healthcare organizations will have to respond to growing cybersecurity challenges, and policymakers will have the responsibility of enacting laws that ensure careful governance and security arrangements for stored data. For example, Google DeepMind's partnership with the Royal Free London NHS (National Health Service) Foundation Trust was severely criticized in 2017 for inappropriate sharing of confidential patient data and their use in an app called Streams that was designed to alert, diagnose and detect acute kidney injury. The Royal Free failed to comply with the UK's Data Protection Act when it handed over the personal data of 1.6 million patients to DeepMind. The ruling of the Information Commissioner's Office (the UK's independent authority set up to uphold information rights in the public interest, promote openness by public bodies and data privacy for individuals) was based largely on the facts that the app continued to undergo testing after patient data were transferred, and that patients were not adequately informed that their data would be used as part of the test (Information Commissioner's Office, undated; Hern, 2017).

Such instances demonstrate the challenges in developing ethical and legal frameworks for data sharing, interoperability of systems, and the ownership of software produced from such partnerships, as well as the legal framework for clinical responsibility when errors occur (*The Lancet*, 2017).

Privacy concerns are also a critical consideration for the use of data. Health data are most often owned by governments, who could be tempted to sell such data on to private companies. In many cases the users can become the 'product' (in effect, patients' data become monetizable). For example, in the US, the Walgreens pharmacy chain collects data contained in prescriptions and sends out mailshots about clinical trials related to the customer's illness. For this service, Walgreens is paid a fee by those recruiting patients for clinical trials and by pharmaceutical companies. Kalev Leetaru, writing in *Forbes* magazine, asserts that: '[…] Walgreens does not explicitly inform customers at purchase time that their prescription may be used to target them for medical trials and offer them the ability to opt-out of having their private medical information used in such a manner […]' (Leetaru, 2018). If companies such as Walgreens are able to do this, then it could be the case that technology companies that gather patient information could also sell individuals' sensitive health data to third parties.

There are further ethical considerations. What obligations do technology companies have to alert populations if their AI produces results that reveal society-wide concerns, such as a potential outbreak of a highly contagious infectious disease? Even if technology companies using AI for health purposes report their findings to governments, history has shown that governments can downplay health risks or fail to alert citizens when economic interests are involved. For example, fears over social and economic stability, as well as the political structure involved in alerting of a disease outbreak, led Chinese leaders to delay reporting the outbreak of Severe Acute Respiratory Syndrome (SARS) in 2003 (Huang, 2004).

Governance is challenging in this realm. Health, technology and data protection policies differ greatly across countries and regions, with many LMIC governments lacking the resources and technological capabilities to create consistent policies on population health. At the same time, many of these countries also lack regulations on the use of data and technology that are intrinsic to AI development.

Accuracy must also be considered. A recent report by the UK Information Commissioner's Office highlights the implications around accuracy of personal data during collection, analysis and application. For example, the results of data analysis may not be representative of the wider population, and hidden biases in datasets can lead to inaccurate predictions about individuals (Information Commissioner's

Office, 2017). Responsibilities are also not clearly defined. Considering the intricate processes involved in AI-produced results, from data collection to algorithm creation and use, how should a government or regulatory system understand who is responsible for flawed AI-derived recommendations?

Algorithms inevitably reflect the bias of training data, and AI tools tend to show a bias reflecting conditions in the high-income countries where they are developed. This is because the algorithms require millions of historical health datapoints, which are often missing in low-resource settings, to provide accurate outputs appropriate to the geography and population (USAID, 2019). Questions about how the AI's algorithms were designed, and with which inputs, remain to be answered as they are central to the questions of their overall utility and of whether they are appropriate for high-, low- and middle-income settings. A recent study by Facebook's AI Lab demonstrates this hidden bias. Five off-the-shelf object recognition algorithms (Microsoft Azure, Clarifai, Google Cloud Vision, Amazon Rekognition, and IBM Watson) were asked to identify household items collected from a global dataset. '[The] object recognition algorithms made around 10 per cent more errors when asked to identify items from a household with a $50 monthly income compared to those from a household making more than $3,500. The absolute difference in accuracy was even greater: the algorithms were 15 to 20 percent better at identifying items from the US compared to items from Somalia and Burkina Faso.' (Vincent, 2019)

Governments – as well as businesses and non-profit organizations developing AI solutions – also need to consider business model sustainability. This will be a challenge in low-resource contexts, where many of the key actors will not have the financial means to purchase these tools.

Health-related AI applications will require strong infrastructural, legal and ethical frameworks. Government-led initiatives to develop and introduce health-related AI applications, across high-, low- and middle-income settings, need to consider these issues. Governments – as well as businesses and non-profit organizations developing AI solutions – also need to consider business model sustainability. This will be a challenge in low-resource contexts, where many of the key actors will not have the financial means to purchase these tools. As one private insurance company representative in East Africa noted: 'I absolutely see the value of AI risk management tools and I realize that this would save us money, but I do not have the budget to buy something now which will save me money 12 months down the line.' (USAID, 2019) This 'applies to many LMIC governments that understand the value of these AI tools, but do not have the resources to buy them, or the human resources or internal IT capabilities to implement them' (ibid.).

Equity issues do not just apply from country to country, but also arise out of the so-called 'digital divide', where different parts of the same society have differing levels of access to advanced technologies such as 4G networks and smartphones. AI tools for health that are enabled by mobile phone technology are only one example of how more connected populations and patients will benefit from services such as medical advice and information through devices to which poorer populations may not have access.

Governments engaging with integrating AI tools into healthcare systems will need to take into consideration not just ethical and legal issues (such as privacy, confidentiality, data security, ownership and informed consent) but also fairness, if AI and related technologies are to contribute to achieving the health-related Sustainable Development Goals (SDG) targets. Ubenwa, an AI application under development in Nigeria, aims to address SDG 3.2 (by 2030, end preventable deaths of newborns and children under five years of age) by providing existing diagnostics that are 95 per cent cheaper than existing clinical software. The AI used is a ML system that can take an infant's cry as input and analyse the amplitude and frequency patterns in the cry to provide an instant diagnosis of birth asphyxia. The test results from Ubenwa's diagnostic software have shown a sensitivity of more than 86 per cent and specificity of 89 per cent. The algorithm has been used in a mobile app that harnesses the processing capabilities of smartphones to provide near-instantaneous assessment of whether or not a newborn has or is at risk of asphyxia (Louise, 2018). Not only is Ubenwa cheaper, and therefore more easily available in low-resource settings; it is also non-invasive (Ubenwa.ai.). Technology trajectories and their impacts will be shaped by local socio-economic contexts, and thus will not be the same everywhere.

India provides a relevant and useful case study to contextualize some of these issues. The government of India recently released its AI strategy, and healthcare is a priority sector for its application in India (Niti Aayog, 2018a). The government seeks to position India as a 'garage' for developing AI solutions for the rest of the world. Many of challenges facing India – from the type of diseases to the quality of the health infrastructure – are shared by a number of other developing economies.

3. AI in Healthcare in India: Applications, Challenges and Risks

There are huge challenges for health systems in India in terms of quality, accessibility, affordability and inequity. On the one hand, India is home to some of the best hospitals in the world, contributing to a growing medical tourism sector (Indo-Asian News Service, 2017). On the other, there is an acute shortage of qualified medical professionals: the ratio of available doctors to population (assuming an availability rate of 80 per cent) can be estimated at 1:1,596 (calculated from Central Bureau of Health Intelligence, 2018). The ratio is particularly low in rural areas, leaving patients to travel long distances to get even basic care. Furthermore, government spending on healthcare is one of the lowest in the world – in 2016–17 only 1.4 per cent of India's GDP was allocated to healthcare (Rao, 2018). Most Indians rely on private health providers – 79 per cent of urban households and 72 per cent of rural households accessed private health facilities in 2014 (National Sample Survey Office, 2014). The private healthcare space, however, is fragmented and unregulated, with approximately 1 per cent of private hospitals in India being formally accredited (Jyoti, 2017). Affordability of healthcare is a further concern: while 30 per cent of total health expenditure is borne by the public sector, patients' out-of-pocket expenses account for the remaining 70 per cent (Rao, 2018). The high cost of private healthcare is a major driver of persistent poverty: in 2011, 55 million Indians were pushed below the official poverty line due to healthcare costs, with 38 million of these falling below the poverty line due to the high cost of medication (Selvaraj, Farooqui and Karan, 2018).

New ML or other AI technologies could help address a number of these challenges, by improving access to quality healthcare, particularly in rural and low-income settings; addressing the uneven ratio of skilled doctors to patients; improving the training and efficiency of doctors and nurses, particularly in complex procedures; and enabling the delivery of personalized healthcare, at scale.

The recently released draft National Strategy for Artificial Intelligence in India highlights that '[the] increased advances in technology, and interest and activity from innovators, [provide an] opportunity for India to solve some of its long existing challenges in providing appropriate healthcare to a large section of its population' (NITI Aayog, 2018a). The government is also trying to create a national digital health infrastructure, as articulated in the recent policy documents for the National Health Stack (2018) (NITI Aayog, 2018b) and the National Digital Health Blueprint (2019) (NDHB) (Ministry of Health and Family Welfare, 2019). Key features of this digital infrastructure include the Healthlocker – an electronic national health registry and cloud-based data storage system that would serve as a single source of health data for the nation; a federated personal health records (PHR) framework that would allow data to be available both to citizens and for medical research; a coverage and claims platform that would support large health protection schemes; a national health analytics platform; and a unique digital health ID for each citizen. The government also launched Ayushman Bharat (Healthy India), or the National Health Protection Scheme (2018), which was devised to provide health insurance to families whose incomes are below the poverty line (India.gov.in, 2018). These build on the earlier National Health Policy (2017), which envisaged creating an integrated health information system linked to the Aadhaar system,[2] and enhancing public health outcomes through big data analytics. These policies call for a state-backed or state-enabled digital infrastructure for data exchange, which is then accessible to the private sector for further innovation, based on open application programming interfaces (APIs) and national data portability (Press Information Bureau, 2019).

The prioritization of AI for healthcare has created an impetus for greater collaboration between government, technology companies and traditional healthcare providers. For example, NITI Aayog, the government's official policy think-tank, is working with Microsoft and the medical technology start-up Forus Health to develop a pilot for early detection of DR.[3] The Maharashtra state government has also signed a memorandum of understanding with NITI Aayog and the Wadhwani AI group[4] to launch the International Centre for Transformational Artificial Intelligence (ICTAI), focusing on rural healthcare (Hebbar, 2018). Similarly, the Telangana state government has adopted the Microsoft Intelligent Network for Eyecare, which was developed in partnership with Hyderabad-based LV Prasad Eye Institute (Gupta, 2018).

Since 2012, about $150 million has been invested in AI start-ups in India, raised predominantly by companies that use AI. Of this, around $77 million was raised in 2017 alone.

Since 2012, about $150 million has been invested in AI start-ups in India (NASSCOM, 2018), raised predominantly by companies that use AI. Of this, around $77 million was raised in 2017 alone. Several new start-ups are already testing and offering a range of solutions automating, for example, analysis of medical tests for the screening and diagnosis of diseases; patient management systems; and early detection and disease prevention systems (Misal, 2018). For example, Niramai, a Bangalore-based start-up, is using ML to detect breast cancer at an early stage. Another start-up, ChironX, employs deep learning algorithms for retinal abnormality detections; and the start-up SigTuple is using AI to create faster diagnostic tools that can enable better primary care and first aid. Large technology companies such as Google and Microsoft have also established research partnerships with leading hospitals to develop and test diagnostic tools (Singh, 2020).

The challenge of delivering quality healthcare at scale presents a strong case for developing AI-based solutions for healthcare in India. However, this process is unlikely to be straightforward or simple, and several questions arise. What are the likely challenges and risks for developing AI-based solutions for healthcare in India? To what extent do they differ from, or resonate with, concerns flagged in global narratives? While AI is still a nascent field in India, with only a small handful of technology companies and start-ups having started to develop and test solutions, early identification of potential risks can help avoid undesirable policy and technological lock-ins. The first part of this chapter identifies some of the key use cases, or areas in which AI is being developed and deployed. It then outlines some of the likely risks and challenges at the different stages of development, adoption and deployment.

[2] Aadhaar is a 12-digit unique identity number based on an individual's biometric and demographic data. See, https://uidai.gov.in.
[3] 3nethra, developed by Forus Health, is a portable device that can screen for common eye problems. Integrating AI capabilities to this device using Microsoft's retinal imaging APIs enables operators of 3nethra devices to obtain AI-powered insights when they are working at eye check-up camps in remote areas with zero or intermittent connectivity to the cloud.
[4] Wadhwani AI is 'an independent nonprofit research institute and global hub, developing AI solutions for social good' (see www.wadhwaniai.org).

Main use cases

Based on desk research and interviews with members of government and industry, we have identified four key areas in which AI solutions for healthcare are being developed. The deployment of AI is still at a very early stage, particularly in the form of clinical interventions. A number of the identified use cases are still at a development and testing stage. Most of the current use cases take the form of decision support systems, followed by process optimization and virtual assistants. Computer vision – one of the more advanced applications of AI – is being used to train AI algorithms to read X-rays and scans to support the processes of disease detection and diagnosis. Only a small handful of companies are developing surgical simulators, personalized health solutions and patient monitoring systems. Moreover, only a small number of interventions use NLP and speech recognition, both of which are critical for meeting diverse linguistic and literacy needs in the country. This is likely to change, however, with growing investments by big tech actors such as Google and Microsoft in the development of these capabilities (*The Times of India*, 2019). Google, Microsoft and IBM have multiple partnerships with private hospital groups such as Narayana, Apollo and Fortis, as well as partnerships with state governments in India. These are working on a range of solutions, including AI systems for hospital management, disease detection and prediction, as well AI service delivery in remote areas (Sinha, 2018).

The four areas in which interventions are being developed are as follows:

Disease detection and diagnostics

ML is being used to build decision-support systems for diagnostics, as well as in predictive systems for prognostication. Computer vision and DL models are being used to read medical scans such as X-rays, CT scans, PET scans and ultrasound scans. AI-based systems are being used for early detection of tumours – e.g. non-invasive, non-touch and non-radiation approaches to detect breast cancer – as well as predicting cancer recurrence through a risk score. AI-based applications are also being developed and used to build systems that can analyse images of blood. SigTuple, for example, is using an AI platform called Manthana for automated analysis of blood smears as well as for the digitization of blood, urine and semen samples (ET Rise, 2018). Researchers at one of India's leading government hospitals have developed a tool that leverages thermal imaging and AI-based tests to help predict the onset of haemodynamic shock. AI systems for tuberculosis diagnosis and DR systems are also being developed. Platforms such as OnliDoc and Lybrate are also using AI methods to provide virtual assistance and diagnostics remotely. OnliDoc uses AI for symptom checking and treatment selection (Misal, 2018).

Process optimization

ML processes are being developed to create new efficiencies in areas such as hospital bed management and processing of insurance claims. A few online platforms that assist with helping find a doctor, storing health records, or procuring medicine are using ML to improve efficiencies in these processes. Others are automating the first-level screening of symptoms, finding doctors and booking appointments. Optical character-reading systems are also being used to scan prescriptions and check prescribed medications against the inventory. ML is also being developed for bed management and planning, to predict rates of 'patient churn' (turnover of beds), in order to optimize the use of beds in hospitals.

Patient-facing applications

Chatbots are increasingly being used as conversational agents for interaction with patients. The online platform mfine, for example, handles more than 15,000 cases per month – approximately the number of patients handled by Manipal Hospitals, one of Bangalore's largest conventional hospital groups. Several large hospitals now use chatbots to schedule appointments, converse, and collect basic details and symptoms, before handing over a case to a doctor. In the case of mental health, chatbots are being used as the first level of intervention in behavioural coaching and as a means of addressing loneliness. Systems to monitor or track patient progress are also under development. AI-driven analysis of camera feeds was found to have been used in one case to detect emotional responses and patient fatigue in order both to help monitor patients during the treatment process and to alert medical staff. Sensor data is also being used to monitor patient recovery and response to medication after surgery or treatment. Wearable sensors and AI-based solutions are being developed to measure vital signs and provide doctors with actionable insights.

Medical R&D and training

While at a much earlier stage of development, DL techniques are being developed to derive molecular insights for drug discovery. Surgery simulators that are continually updated are also being developed to train doctors for spine and knee surgery. A surgery simulator centre was recently opened in Delhi.

The main sources of data for developing these systems are primarily historical data held within research institutions, non-profit organizations and medical service providers. In some cases, data are collected by developers through other platforms or healthcare services they already provide. For example, a healthcare platform that enables doctor discovery, online consultations and online medical purchase would then employ user data captured on their platform to build an AI system to optimize and automate certain operations pipelines – doctor-to-patient matching, and doctor discovery based on location data – in their services.

> The main sources of data for developing these systems are primarily historical data held within research institutions, non-profit organizations and medical service providers. In some cases, data are collected by developers through other platforms or healthcare services they already provide.

In cases where the models need to be trained on more general aspects, such as conversational ability or image recognition, developers use open-source data to get the model off the ground. However, as India does not have robust medical datasets, start-ups often use publicly available datasets from the US and Europe. If the AI algorithm is a small part of their systems, or they do not possess in-house capabilities, developers also use models available with cloud providers such as Google, Microsoft and Amazon.[5] In certain cases, new data are being collected through field experiments; one such example is Wadhwani AI, which is seeking to build a model to estimate and document the weight of a child at birth for a public health census (Goyal, 2019).

[5] Insights from fieldwork conducted by Tandem Policy Lab, May 2019.

Additionally, a few start-ups and initiatives have begun to provide personalized health solutions. Healthi, a digital health and wellness start-up in Bangalore, uses predictive analytics, personalization algorithms and ML to deliver personalized health suggestions. Similarly, Manipal Hospitals is using IBM Watson for Oncology, a cognitive-computing platform, to help physicians discover personalized cancer care options.

Challenges and risks

While AI could bring benefits to healthcare in India, it will certainly not provide a simple solution or fix in an already very complex health landscape involving numerous stakeholders, competing priorities, entrenched incentive systems and institutional cultures. Further new and complex challenges will also arise around data use, privacy and security. This section explores the challenges and risks involved in using AI systems for healthcare solutions in India across three stages: development, adoption and deployment.

Development

While a number of factors are relevant at development stage, two are of particular note: the challenges entailed in obtaining structured, complete and representative datasets; and the human, financial and infrastructural resources needed to build AI solutions. While these challenges exist globally, they are further accentuated in the Indian context.

Access to data

AI systems depend on the availability of large amounts of data. This poses a major impediment for building indigenous AI interventions in India. Datasets for healthcare in India are fragmented, dispersed and incomplete.[6] Building AI requires longitudinal data. People often go to different doctors, even for the same diagnosis or treatment; even large hospitals do not have loyal patient followings. A large proportion of these healthcare providers are unaccredited and informal health practitioners, with non-standardized data collection, recording and analysis systems, and differing approaches to medical care more generally.

Digitization practices are poor, uneven and not standardized; and there is no centralized database for health records in India. Even in large hospitals, it is often the case that every time a patient visits, a new registration number or patient file is created for them, and doctors' prescriptions are handwritten. Frontline health workers in India record patient histories in notebooks, using their own systems of annotation. The data that are readily available for AI companies are thus likely to be unrepresentative of a significant part of the population.

Health policy is also determined at a state level, and not by the central government. This means there are significant variations across states as well as differing levels of readiness to share health data. State governments are often reluctant to share population data because this may reflect poorly on their capacities for governance; misrepresentation or fudging of health data is also not uncommon.[7] Inadequate or unrepresentative data can result in poor data quality and coherence, leading to erroneous algorithms and possible misdiagnosis.

[6] Insights from Tandem Policy Lab, 3 and 4 December 2018.
[7] Insights from Tandem Policy Lab, 3 and 4 December 2018.

Efforts to digitize the health system are now under way. Plans for an Integrated Health Information Program (IHIP) to create EHRs for all citizens, and enable the interoperability of existing EHRs, are currently in development (National Health Portal of India, 2017). However, implementation of EHRs is not harmonized, leading to different interpretations of record digitization and data retention (Paul et al., 2018). Moreover, health workers are often overworked, and may be unable or unwilling to invest the time and effort needed. Capacity limitations have also led to time lags between a health episode and its digitization.

The government is also attempting to create a digital health infrastructure that can enable AI solutions. A recent discussion paper by NITI Aayog outlined ambitions for creating a 'National Health Stack (NHS)' to make both personal health records and service provider records available on cloud-based services to private healthcare actors. The NHS is expected to consist of four key elements – electronic health registries of health service providers and beneficiaries, a coverage and claims platform, a federated personal health records framework and a national health analytics platform (NITI Aayog, 2018b).

While having such open data stacks could enable private-sector innovation, a number of issues are yet to be examined by health and other relevant ministries. As noted earlier, datasets are incomplete and unrepresentative, and are scattered across thousands of healthcare providers. Making these data machine readable will be an enormous undertaking, requiring not only human and financial resources, but also coordination across an enormous range of healthcare providers. Moreover, the business case for healthcare providers to share their data has not been examined. The consent-based model that is being proposed as central to India's data protection framework is also likely to be inadequate. As Mayer-Schönberger and Cukier argue: 'In the era of big data, the three core strategies long used to ensure privacy – individual notice and consent, opting out, and anonymization – have lost much of their effectiveness' (Mayer-Schönberger and Cukier, 2017). Beyond the issue of privacy, fundamental questions remain unaddressed as to who owns healthcare data, who should be allowed to use it, and in what way.

Blind spots in data collection

Current AI experiments are dependent on historical data available from select hospitals or research institutes. The trouble with historical data, as has already been well documented in existing studies on AI, is that it will, by definition, reflect certain societal structures of discrimination (Gershgorn, 2018). For example, there are already documented instances of women from lower castes being denied healthcare due to the medical provider's class elitism (Siddiqui, 2008). Similarly, data from clinical trials being used to inform AI typically under-represent women, minorities and the elderly, as fewer of them are selected for such trials (Hart, 2017). As a result, the medicines formulated through these data are effective only for certain populations. Algorithms trained on these datasets thus risk having certain blind spots, which could in extreme cases lead to misdiagnosis.

In other cases, start-ups are also working with open data repositories, but much of this is for populations in other geographies. This could result in algorithms that are not easily applicable for Indian populations, again contributing to a risk of misdiagnosis. For instance, India's Manipal Hospitals has linked with IBM Watson for Oncology to aid doctors in the diagnosis and treatment of seven types of cancer. But a number of physicians have already noted that the population on which Watson is trained does not accurately reflect the diversity of cancer patients across the world, and, as a result, the system is heavily biased towards US patients and standards of care (Ross and Swetlitz, 2017).

Infrastructure and costs
AI systems can be expensive to train, test and deploy. Datasets are expensive to collect, and computing power and storage space is expensive. Most healthcare organizations also lack the data infrastructure necessary to collect the data needed to optimally train algorithms – i.e. to test them for bias and adjust the model, and continually monitor and evaluate field outcomes. The unavailability of digital infrastructure required to build AI systems is a further constraint. Cloud-based computing infrastructure is mostly concentrated in servers outside India. As a result, many start-ups have also incorporated themselves outside India. Moreover, as the commentator Shashi Shekhar Vempati noted in a 2016 report, the lack of technological infrastructure has made it difficult to develop applications based on DL techniques. This poses a major challenge for developing AI capacities across different languages, which would be particularly relevant for the adoption of AI in primary health across diverse rural contexts (Vempati, 2016). The shortage of skilled data scientists is a further impediment.

> Most Indian healthcare companies do not have the computing power to build AI systems, and rely on cloud services provided by the big tech players – Google and Microsoft in particular.

In practice, these infrastructural constraints frequently lead both Indian hospitals and start-ups that are keen to leverage AI for healthcare to a dependence on a few very large technology companies such as Google and Microsoft. In the case of the Forus-Microsoft partnership for DR screening, the device, built by Forus, sends an image of the screened patient to the cloud, following which the algorithm screens it for DR and sends its interpretation back to the device or the doctor (Singh, 2020). Large technology companies, which already have an advantage because of the large quantity of user data they possess, are likely to have a further advantage over smaller health actors. Lack of regulation also makes data acquisition easier for big tech firms. For example, according to Seema Singh, Google had previously attempted to obtain data from US medical establishments for developing AI solutions for eyecare, but was unable to gain approval beyond what was already in the public domain. It then turned to India, and has now struck up numerous data sharing collaborations with eye hospitals in India (ibid.). Most Indian healthcare companies do not have the computing power to build AI systems, and rely on cloud services provided by the big tech players – Google and Microsoft in particular. In fact, as Singh has commented, big tech is primarily in the healthcare business in order to benefit the sale of its cloud software (ibid.).

Adoption

While adoption is likely to be shaped by a number of factors, two in particular are likely to be particularly relevant in India: the infrastructural and financial feasibility of adoption; and the degree of readiness and acceptance within established healthcare practices.

Affordability and infrastructure
Much of the dominant narrative around AI for healthcare in India focuses on the potential to reach underserved populations, particularly in rural areas lacking infrastructure or sufficient physicians, or among economically weaker sections of society where the population lacks the financial means to access medical facilities (Paul et al., 2018).

However, a closer look at AI adoption reveals that, while a number of pilots are being run for rural healthcare, the bulk of adoption is limited to large private hospitals or clinics. For example, the IBM Watson platform for cancer diagnostics was first implemented by the Manipal and Apollo hospital groups, along with other private hospitals (Kambli, 2019). This is likely to be for reasons of affordability: private medical practitioners, smaller clinics and rural hospitals are unlikely to have the financial means to adopt these solutions or to have spare resources for piloting and experimentation.[8] In diagnostics, for example, it would be safe to assume that AI-based solutions are going to be adopted by institutions that already have diagnostic capabilities, such as diagnostic labs and larger hospitals. Some technology companies are exploring partnerships with existing medical device manufacturers to integrate AI solutions into existing devices. However, these solutions are likely to be unviable in the many rural settings that have poor internet connectivity and little digital infrastructure.[9] In a select few hospitals in urban settings, these AI solutions are being advertised to patients as an additional layer of diagnosis available to those able to pay the additional costs, thus risking an exacerbation of existing inequities within health systems in India.

This raises important questions about the distribution of AI gains; for instance, as to whether these gains will be limited to the elite few. In contrast to the promise of improving access to quality healthcare for underserved populations, there is a risk that affordability and the lack of a sustainable business case for AI in rural healthcare could be severe constraints for adoption. At present, the biggest winners in the Indian context seem to be the big tech companies who between them provide the cloud infrastructure for most hospitals. Big tech has signed up to numerous partnerships with private hospitals, but the terms of such agreements are not publicly available. Questions around who owns the data, and how value is going to be generated and shared across various stakeholders, are yet to be addressed (Singh, 2020).

> In contrast to the promise of improving access to quality healthcare for underserved populations, there is a risk that affordability and the lack of a sustainable business case for AI in rural healthcare could be severe constraints for adoption.

Another critical issue is the process for obtaining medical approvals for adoption and deployment. In the US, the Food and Drug Administration has only recently proposed a framework for dealing with AI in medical devices. In the Indian context, where there are low levels of institutional capacity and acute healthcare needs, big tech has been able to circumvent some of these regulatory hurdles. For example, in a press release in August 2019, Microsoft claimed it had 'screened' more than 200,000 people for cardiovascular diseases using 'the AI-powered API across Apollo Hospitals', even predicting the risk score for some. However, there is no peer review publication yet. In another example, Google recently published a study where its DL models looked at around 600,000 chest X-rays from Apollo hospitals (Majkowska, Mittel et al., 2019). However, the same dataset was used for training and testing the data – a bad practice that is none the less widely established, and which can be used to show high success rates that are not necessarily valid (Singh, 2020).

[8] Insights from Tandem Policy Lab, 3 and 4 December 2018.
[9] Insights from fieldwork conducted by Tandem Research, 10 May–22 July 2019.

Institutional capacities and cultural acceptance

Adoption in rural or underserved areas is likely to depend on existing levels of government support, rather than on direct access to populations or market solutions. As noted earlier, healthcare is state-run in India. Individual states have differing institutional capacities and knowledge systems. Many states may not have the technical expertise to oversee health regulation or develop an ecosystem that encourages innovation and improves access. Furthermore, there are often instances of data fudging among low-performing states and a general reluctance to make data available for public scrutiny. This could also mean that low-performing states are hesitant to permit external scrutiny and intervention.[10]

Moreover, AI innovations will not by themselves change the incentives that support existing ways of working in the healthcare sector (Rajkomar, Dean and Kohane, 2019). A complex web of ingrained political and economic factors, along with medical practice norms and commercial interests, determine the way healthcare is delivered. Simply adding AI applications to a fragmented system will not create sustainable change (ibid.). The healthcare sector has also traditionally been resistant to the permeation of information and communications technologies (ICT) (Safi, Thiessen and Schmailzl, 2018). Medical professionals often find it time-consuming and laborious to change their standard way of working; some also see it as a form of management control.[11] In many cases, uptake has been mostly symbolic, to satisfy management or reporting requirements. Doctors still rely on and prefer handwritten files; in some cases, even where patient data are entered into a digital database, the electronic record is deleted after a printout is taken and filed away (Powell, Tyagi and Ludhar, 2018). Cultural and social attitudes are likely to shape the speed and scale of adoption. For example, a recent study found that women in rural areas tend to seek out informal healers over formal healthcare providers; this was related to factors such as ease of communication, cultural familiarity or resonance, avoidance of social stigma and geographical distance from formal healthcare facilities (Das et al., 2018).

Deployment

Three key sets of challenges will need to be considered at the level of deployment: privacy, misuse and accountability.

Privacy

Healthcare data are highly sensitive, and data breaches can have implications for an individual's personal autonomy, dignity, and even access to work. In 2016, the hacking of a Mumbai-based diagnostic laboratory database led to the leaking of medical records (including HIV status reports) of more than 35,000 patients. This database held the records of patients across India, and many may still be unaware that their details have been exposed. The database had been subjected to multiple hacks in the previous few years, sometimes up to three times a week. However, no action had been taken by the laboratory concerned to secure the data (Express News Service, 2016).

In March 2018 India's Ministry of Health and Family Welfare released to the public domain a draft of the proposed Digital Information Security in Healthcare Act (DISHA 2018), which would enable the digital sharing of personal health records between hospitals and clinics. DISHA would provide for a rights-based framework for medical privacy, conferring on patients the right to privacy, confidentiality and security. The draft law requires that each instance of transmission

[10] Insights from Tandem Policy Lab, 3 and 4 December 2018.
[11] Insights from Tandem Policy Lab, 3 and 4 December 2018.

of digital health data gains the explicit prior permission of the owner, and patients would have the right to refuse consent for the generation, storage and collection of their data (Ministry of Health and Family Welfare, 2017).

However, such a framework gives rise to a number of challenges. First, obtaining meaningful consent would require the entire population to have the capacity to make informed decisions about the collection and use of their data. In many cases, patients have low levels of literacy and education; in other cases, there may be consent fatigue, particularly where terms and conditions are difficult and time-consuming to comprehend (Bailey et al., 2018).

The draft legislation also does not contain provisions for instances where the digital health data of the owner have been collected without his or her consent: neither does it mention the status of the data when the owner withdraws their consent (Mohandas and Sinha, 2018). Furthermore, in an AI-equipped world, anonymization of data is not enough: recent studies show that triangulation across multiple data points makes it possible to identify individual users (Culnane, 2017). In addition, with ML models, it is difficult – if not impossible – to identify how a particular piece of data is being interpreted and used for building algorithmic models; it is thus computationally very difficult to ascertain how a particular data source is being used, or to withdraw consent for its use (Shou, 2019).

The current draft Personal Data Protection (PDP) bill, yet to be passed by the Indian parliament, also proposes a consent-based model. The PDP draft legislation designates health data to be 'critical' and 'sensitive', requiring a set of permissions from the owner before being used. Given that there are three actors in health data co-creation – the patient, the hospital or doctor, and the payer – the bill ignores consent and rights at different stages of the data life cycle. Consent-based frameworks are also inexact as to the amount of patient data that will be disclosed to private players in the system such as insurers, pharmacies and hospitals.

> Further concerns arise on account of the linking of health data with the controversial biometric identity project Aadhaar, which has already been documented as suffering multiple privacy and security breaches.

Further concerns arise on account of the linking of health data with the controversial biometric identity project Aadhaar, which has already been documented as suffering multiple privacy and security breaches (Business Line, 2018). There is no clear understanding of how the Aadhaar data will be used, and who will have access to them.

There also seems to be a dissonance between existing digitization policies and privacy policies for healthcare. The draft DISHA prohibits the use of digital health data for commercial purposes, especially by insurance companies, employers, human resource consultants and pharmaceutical companies. However, the proposed NHS has a special platform dedicated to insurance claims and coverage. It remains to be seen how the data protection provisions of the DISHA are going to apply to the NHS. Current policy frameworks seem to be torn between the need to promote innovation and AI development, and the need to have the right frameworks to protect user privacy and establish user control.

The issue of consent must also be considered in the context of the nature of the doctor-patient relationship. As Karunakaran Mathiharan noted in a 2014 report, doctor-patient relationships are characterized by power differentials and cultural notions in India, where a doctor's authority is often considered absolute and where they are accorded a high level of trust from patients. This is especially true since a significant part of the population falls outside the ambit of formally recognized medical systems, rendering moot the issue of obtaining informed consent (Mathiharan, 2014).

Misuse

The linking of health data with other systems, and the new avenues for discrimination this may create, gives rise to significant concern. Health insurance data, for example, can be leveraged by banks to evaluate eligibility for loans: a poor performance on health indicators could be seen as an indication of an individual's inability to work, which would increase the likelihood of non-payment. The flow of health data to companies outside the healthcare sector could lead to discrimination in the workplace or in other entitlement and social benefits.

Issues around data security are of equal importance. The Aadhar system, for example, has already been subject to multiple data breaches, with only weak attempts having been made to improve the security infrastructure (Vidyut, 2018). The Digilocker – on which the Healthlocker would be modelled – also has inadequate security measures, raising concerns about the biometric data stored within it. Security concerns are also entailed in the initiation by the government of India of the eSign electronic signature framework (which allows an Aadhaar cardholder to digitally sign a document), since third parties are involved (Jalan, 2019). Sensitive health data will become vulnerable if this model is followed. Although the NDHB refers to the creation of a Security Operations Centre (SOC) and a NDHB Security Policy, it crucially does not cite the procedures to be followed in case of a privacy breach (ibid.). Cyberattacks on medical institutions can be used to tamper with the data or create fake health records. Once cyberattackers have access to an institution's systems and health records, they can encrypt all the systems to be completely inaccessible and unusable by the victimized medical institutions and demand a ransom. Vulnerability to cyberattacks arises due to many factors: outdated digital infrastructure, for instance, or a lack of awareness and training among medical personnel on the subject of cyberattacks (NovoJuris Legal, 2019). As reported by multiple news agencies, in June 2018 Mahatma Gandhi Memorial Hospital, a trust-run hospital in Mumbai, was affected by a ransomware attack in which hospital administrators found their systems to be locked, subsequently noticing an encrypted message sent by the attackers demanding a ransom payment in bitcoin in exchange for the unlocking of the system. It was reported that the hospital had lost 15 days' data related to billing and patients' histories, although the hospital did not incur any additional financial loss (Purandare, 2018).

Accountability

Finally, there is the question of accountability. Who is to be held accountable in the case of misdiagnosis or error? On the one hand, AI systems are currently being envisaged as decision-support systems. They are intended not to replace doctors, but to provide a first layer of screening. In other words, the expectation is that there will be a 'human in the loop' to interpret results and point out any errors. However, it is worth asking what type of professional this human might be; and what might be their capacities, and incentives, to check the validity of suggestions produced by an AI system. In rural settings, for example, front-line health workers may not have the knowledge, training or confidence to be able to interpret and challenge AI-generated results. In contexts where doctors are overwhelmed by the number of patients they are treating, and are under pressure to demonstrate efficiency, they may not have the time or incentive structures to correct AI systems in the event of an error. Over-reliance on decision-support systems can create complacency, leading to errors, be they due to blind adherence or inaction (Wickens et al., 2015). These concerns are further accentuated in the context of the weak regulation of the Indian health sector. In recent years, there have been numerous reports of negligence and malpractice of even well-established private hospitals. The main reason that these violations are all too common is the lack of strict and uniform regulation of healthcare in the country (Narayanan, 2017).

Moreover, the question of accountability becomes more complicated when individual health data are being used to aid other decision-making processes – such as credits or loan applications – particularly since the ways in which predictive and self-learning algorithms draw inferences or patterns are hard to identify. Within the AI programmer community, there is a movement to explore fairness, accountability and transparency (FAT) frameworks. But, as others have noted, fairness is a property of social relations, not of code (Selbst et al., 2019). The algorithms must be audited, not only for efficiency and accuracy, but also for issues of social context such as biases and knock-on effects. For instance, the implications of the use of an AI algorithm by a doctor, as opposed to a front-line community health worker (such as an accredited social health activist – ASHA), could be very different, simply due to each practitioner's individual level of training andability to be critical of the output.

The use of AI is likely to transform patient-doctor relations, and related systems and rituals of trust. Much of what transpires between doctors and patients relates to relationship-building, and not merely the provision of medical expertise. We then need to ask what safeguards are needed to build trust and encourage buy-in; and how patients are brought into the processes of deliberation and explanation.

4. Conclusion: Considerations for AI and Healthcare in India

AI can undoubtedly bring new efficiencies and quality to healthcare outcomes in India. However, gaps and challenges in the healthcare sector reflect deep-rooted issues around inadequate funding, weak regulation, insufficient healthcare infrastructure, and deeply embedded socio-cultural practices. These cannot be addressed by AI solutions alone.

Moreover, technological possibility cannot be equated to adoption. In India, poor digital infrastructure, a large, diverse and unregulated private sector, and variable capacity among states and medical professionals alike, mean that the adoption of AI is likely to be slow and deeply heterogeneous. The same factors also make it quite likely that well-established private hospitals will be the main adopters. This in turn would imply that much of the dominant narrative or rationale for the development of AI in healthcare, in terms of improving equity and quality, is unlikely to be addressed through market forces alone: these solutions are more likely to serve populations who already have access to high-quality care, typically in cities with well-developed digital infrastructure. In many small hospitals and single-provider practices in India, administrative systems have barely moved beyond rudimentary ICT solutions such as invoicing and billing platforms.

The effectiveness of these systems will depend on accurate identification of problems and their matching to appropriate solutions. Currently, there is a risk that solutions are technology-led rather than problem-led, and they are as a result often blind to specific contextual needs or constraints. For example, it might not be the best approach to design real-time or synchronous solutions for digital products meant to be used in remote areas where basic internet infrastructure is lacking. Designing the right digital interventions is often challenging because of the digital divide between the user and the technology developers, who are typically more adept at using technology than the user is (Deo and Tyagi, 2019). Finally, issues around privacy, misuse and accountability are only slowly being understood, and require much more far-reaching consideration before AI can deliver safe and fair healthcare solutions.

References

Accenture (2017), 'Artificial Intelligence: Healthcare's New Nervous System', https://www.accenture.com/_acnmedia/PDF-49/Accenture-Health-Artificial-Intelligence.pdf#zoom=50 (accessed 5 Mar. 2020).

Ash, M., Petro, J. and Rab, S. (2019), 'How AI in the Exam Room Could Reduce Physician Burnout', *Harvard Business Review*, 12 November 2019, https://hbr.org/2019/11/how-ai-in-the-exam-room-could-reduce-physician-burnout (accessed 5 Mar. 2020).

Bailey, R., Parsheera, S., Rahman, F. and Sane, R. (2018), *Disclosures in privacy policies: Does "notice and consent" work?*, Working Paper, No. 246, New Delhi: National Institute of Public Finance and Policy, https://www.nipfp.org.in/media/medialibrary/2018/12/WP_246.pdf (accessed 10 Mar. 2020).

Business Line (2018), '1 bn records compromised in Aadhaar breach since January: Gemalto', *The Hindu Business Line*, 20 October 2018, https://www.thehindubusinessline.com/news/1-bn-records-compromised-in-aadhaar-breach-since-january-gemalto/article25224758.ece (accessed 11 Mar. 2020).

Central Bureau of Health Intelligence (2018), 'National Health Profile 2018', Ministry of Health and Family Welfare, http://www.cbhidghs.nic.in/WriteReadData/l892s/Before%20Chapter1.pdf (accessed 10 Mar. 2020).

Colclough, G., Dorling, G., Riahi, F., Ghafur, S. and Sheikh, A. (2018), *Harnessing Data Science and AI in Healthcare: From policy to practice*, Report of the WISH Data Science and AI Forum 2018, https://www.wish.org.qa/wp-content/uploads/2018/11/IMPJ6078-WISH-2018-Data-Science-181015.pdf (accessed 30 Mar. 2020).

Culnane, C., Rubinstein, B. I. P. and Teague, V. (2017), 'Health Data in an Open World', *arXiv preprint,* 15 December 2017, arXiv:1712.05627 (accessed 10 Mar. 2020).

Dargan, R (2019), 'AI Applications Aid Efficiency, Ease Workload', Radiological Society of North America, 16 July 2019, https://www.rsna.org/news/2019/july/ai-applications-address-workload (accessed 5 Mar. 2020).

Das, M., Angeli, F., Krumeich, A. J. S. M., and van Schayck, O. C. P. (2018), 'The gendered experience with respect to health-seeking behaviour in an urban slum of Kolkata, India', *International Journal for Equity in Health*, 17(1): p. 24. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5813424/#__ffn_sectitle (accessed 11 Mar. 2020).

Davenport, T. and Kalakota, R. (2019), 'The potential for artificial intelligence in healthcare', *Future Healthcare Journal,* 6(2): pp. 94–98, doi:10.7861/futurehosp.6-2-94 (accessed 19 May, 2020).

Deo, S. and Tyagi, H. (2019), 'Rx ICT: Digitally disrupting healthcare', *Forbes India*, 22 February 2019, http://www.forbesindia.com/article/isbinsight/rx-ict-digitally-disrupting-healthcare/52629/1 (accessed 10 Mar. 2020).

ET Rise (2018), 'How Manthana uses AI to process visual medical data', *Economic Times*, 17 August 2018, https://economictimes.indiatimes.com/small-biz/startups/newsbuzz/how-manthana-uses-ai-to-process-visual-medical-data/articleshow/65443587.cms (accessed 11 Mar. 2020).

Expert System (2017), 'What is Machine Learning? A definition', 7 March 2017, https://www.expertsystem.com/machine-learning-definition (accessed 5 Mar. 2020).

Express News Service (2016), 'Maharashtra website hacked: Diagnostic lab details of 35,000 patients leaked', *The Indian Express*, 3 December 2016, https://indianexpress.com/article/india/diagnostic-lab-details-of-35000-patients-leaked-hiv-reports-4407762 (accessed 10. Mar 2020).

Forbes (2019), 'AI and Healthcare – a Giant Opportunity', 11 February 2019, https://www.forbes.com/sites/insights-intelai/2019/02/11/ai-and-healthcare-a-giant-opportunity/#60febfc4c682 (accessed 5 Mar. 2020).

Garrett, P. and Seidman, J. (2011), 'EMR vs EHR – What is the Difference?', Health IT Buzz, 4 January 2011, https://www.healthit.gov/buzz-blog/electronic-health-and-medical-records/emr-vs-ehr-difference (accessed 19 May 2020).

Gershgorn, D. (2018), 'If AI is going to be the world's doctor, it needs better textbooks', *Quartz*, 6 September 2018, https://qz.com/1367177/if-ai-is-going-to-be-the-worlds-doctor-it-needs-better-textbooks (accessed 11 Mar. 2020).

Goyal, M. (2019), 'How Wadhwani brothers Sunil and Romesh are using AI to serve the underserved', *The Economic Times*, 5 March 2019, https://economictimes.indiatimes.com/tech/software/how-wadhwani-brothers-sunil-and-romesh-are-using-ai-to-serve-the-underserved/articleshow/68268245.cms (accessed 11 Mar. 2020).

Gupta, P. (2018), 'Bringing together an AI network for eyecare to prevent avoidable visual impairment', *The Microsoft India Blog*, 14 March 2018, https://docs.microsoft.com/en-us/archive/blogs/msind/mine-ai-network-for-eyecare (accessed 11 Mar. 2020).

Hart, D. R. (2017), 'If you're not a white male, artificial intelligence's use in healthcare could be dangerous', *Quartz*, 10 July 2017, https://qz.com/1023448/if-youre-not-a-white-male-artificial-intelligences-use-in-healthcare-could-be-dangerous (accessed 11 Mar. 2020).

Haskew, J., Rø, G., Saito, K., Turner, K., Odhiambo, G., Wamae, A., Sharif, S. and Sugishita, T. (2015), 'Implementation of a cloud-based electronic medical record for maternal and child health in rural Kenya', *International Journal of Medical Informatics*, 84(5), pp. 349–54, doi:10.1016/j.ijmedinf.2015.01.005 (accessed 30 Mar. 2020).

Health IT Analytics (undated), 'Google Pursues AI, Voice Recognition for EHR Workflows', https://healthitanalytics.com/news/google-pursues-ai-voice-recognition-for-ehr-workflows (accessed 5 Mar. 2020).

Health Nucleus (undated), 'Artificial Intelligence & Machine Learning'..

Hebbar, P. (2018), 'Maharashtra CM Launches New Artificial Intelligence Centre For Rural Healthcare', *Analytics India Magazine,* 14 November 2018, https://analyticsindiamag.com/maharashtra-cm-launches-new-artificial-intelligence-centre-for-rural-healthcare (accessed 16 Apr. 2020).

Hern, A. (2017), 'Royal Free breached UK data law in 1.6m patient deal with Google's DeepMind', *Guardian*, 3 July 2017, https://www.theguardian.com/technology/2017/jul/03/google-deepmind-16m-patient-royal-free-deal-data-protection-act (accessed 5 Mar. 2020).

Huang, Y. (2004), 'The SARS Epidemic and its Aftermath in China: a Political Perspective', in Institute of Medicine (US) Forum on Microbial Threats, Knobler, S., Mahmoud, A., Lemon, S. et al. (eds) (2004), *Learning from SARS: Preparing for the Next Disease Outbreak: Workshop Summary*, Washington, DC: National Academies Press, available at: https://www.ncbi.nlm.nih.gov/books/NBK92479 (accessed 16 Mar. 2020).

Information Commissioner's Office (undated), https://ico.org.uk/ (accessed 30 Mar. 2020).

Information Commissioner's Office (2017), 'Big data, artificial intelligence, machine learning and data protection,' https://ico.org.uk/media/for-organisations/documents/2013559/big-data-ai-ml-and-data-protection.pdf (accessed 30 Mar. 2020).

India.gov.in (2018), 'Ayushman Bharat: National Health Protection Mission', https://www.india.gov.in/spotlight/ayushman-bharat-national-health-protection-mission (accessed 10 Mar. 2020).

Indo-Asian News Service (2017), 'India is now among fastest-growing medical tourism destinations', *India Today*, 13 October 2017, https://www.indiatoday.in/travel/india/story/india-medical-tourism-asia-fastest-growing-yoga-ayurveda-advantage-healthcare-india-summit-medical-treatment-karnataka-ficci-report-us-russia-saudi-arabia-united-arab-emirates-lifest-1064130-2017-10-13 (accessed 11 Mar. 2020).

Jalan, T. (2019), 'Centre for Internet & Society's comments on NDHB: Literacy and benefits of PHR, access to anonymised data, consent as product', *Medianama*, 13 August 2019, https://www.medianama.com/2019/08/223-centre-for-internet-societys-comments-on-ndhb (accessed 11 Mar. 2020).

Joshi, N. (2019), 'How Far Are We From Achieving Artificial General Intelligence?', *Forbes*, 10 June 2019, https://www.forbes.com/sites/cognitiveworld/2019/06/10/how-far-are-we-from-achieving-artificial-general-intelligence/#8f678606dc4d (accessed 2 Feb. 2020).

Jyoti, A. (2017), '1 per cent of India's hospitals apply for NABH accreditation', *The Pioneer*, 2 November 2017, https://www.dailypioneer.com/2017/india/1-per-cent-of-indias-hospitals--apply-for-nabh-accreditation.html (accessed 10 Mar. 2020).

Kambli, R. (2019), 'Watson Health has more than 15,000 clients and partners', *Express Healthcare*, 5 June 2019, https://www.expresshealthcare.in/healthcare-it/watson-health-has-more-than-15000-clients-and-partners/411517 (accessed 11 Mar. 2020).

The Lancet (2017), 'Artificial intelligence in health care: within touching distance', *The Lancet*, 390(10114), doi:10.1016/S0140-6736(17)31540-4 (accessed 30 Mar. 2020).

Leetaru, K. (2018), 'How Data Brokers and Pharmacies Commercialize our Medical Data', *Forbes*, 2 April 2018, https://www.forbes.com/sites/kalevleetaru/2018/04/02/how-data-brokers-and-pharmacies-commercialize-our-medical-data (accessed 16 Mar. 2020).

Lehne, M., Sass, J., Essenwanger, A., Schepers, J. and Thun, S. (2019), 'Why digital medicine depends on interoperability', *npj Digital Medicine*, 2, article number 79, https://www.nature.com/articles/s41746-019-0158-1 (accessed 19 May 2020).

Louise, N. (2018), 'This Nigerian artificial intelligence health startup is saving thousands of babies' lives with a simple app', Techstartups.com, https://techstartups.com/2018/01/14/this-nigerian-artificial-intelligence-health-startup-is-saving-thousands-of-babies-lives-with-a-simple-app (accessed 5 May 2020).

Majkowska, A., Mittal, S., Steiner, D. F., Reicher, J. J., McKinney, S. M., Duggan, G. E., Eswaran, K., Chen, P. C., Liu, Y., Kalidindi, S. R., Ding, A., Corrado, G. S., Tse, D. and Shetty, S. (2019), ' Chest Radiograph Interpretation with Deep Learning Models: Assessment with Radiologist-adjudicated Reference Standards and Population-adjusted Evaluation', *Radiology*, 294(2): pp. 421–31, doi:10.1148/radiol.2019191293 (accessed 10 Mar. 2020).

Mathiharan, K. (2014), 'Law on consent and confidentiality in India: A need for clarity', *National Medical Journal of India*, 27: pp. 39–42, http://archive.nmji.in/archives/Volume-27/Issue-1/27-1-SFM-III.pdf (accessed 10 Mar. 2020).

Mayer-Schönberger, V. and Cukier, K. (2017), *Big Data: the Essential Guide to Work, Life and Learning in the Age of Insight*, London: John Murray.

Ministry of Health and Family Welfare (2017), *Digital Information Security in Healthcare Act,* Government of India, November 2017, https://www.nhp.gov.in/NHPfiles/R_4179_1521627488625_0.pdf (accessed 16 Apr. 2020).

Ministry of Health and Family Welfare (2019), *National Digital Health Blueprint*, April 2019, https://www.nhp.gov.in/NHPfiles/National_Digital_Health_Blueprint_Report_comments_invited.pdf (accessed 16 Apr. 2020).

Misal, D. (2018), '11 Indian Startups Revolutionising The Healthcare Sector With AI', *Analytics India Magazine*, 26 October 2018, https://www.analyticsindiamag.com/11-indian-startups-revolutionising-the-healthcare-sector-with-ai (accessed 11 Mar. 2020).

Mohandas, S. and Sinha, A. (2018), 'Comments on the Draft Digital Information Security in Healthcare Act', The Centre for Internet and Society, 21 April 2018, https://cis-india.org/internet-governance/files/comments-on-draft-digital-information-security-in-healthcare-act (accessed 11 Mar. 2020).

Narayanan, N. (2017), 'Interview: "We need a medical regulator that is more like a tribunal and not the consumer court"', *Scroll.in*, 11 December 2017, https://scroll.in/pulse/861096/interview-we-need-a-medical-regulator-that-is-more-like-a-tribunal-and-not-the-consumer-court (accessed 10 Mar. 2020).

Narayanan, N. (2018), 'Niti Aayog plan for Aadhaar-linked digital health records raises concerns over safety and privacy', *Scroll.in*, 12 July 2018, https://scroll.in/pulse/886153/niti-aayog-plan-for-aadhaar-linked-digital-health-records-raises-concerns-over-safety-and-privacy (accessed 11 Mar. 2020).

NASSCOM (2018), 'NASSCOM Artificial Intelligence Primer 2018', https://community.nasscom.in/wp-content/uploads/attachment/nasscom-ai-primer-2018.pdf (accessed 16 Apr. 2020).

National Health Portal of India (2017), 'National eHealth Authority (NeHA)', *e-Health India*, 29 May 2017, https://www.nhp.gov.in/national_eHealth_authority_neha_mtl (accessed 11 Mar. 2020).

National Sample Survey Office (2014), 'Health in India: NSS 71st Round (January–June 2014)', Government of India, http://mospi.nic.in/sites/default/files/publication_reports/nss_rep574.pdf (accessed 10 Mar. 2020).

Neelakantan, M. (2018), 'NITI Aayog's National Health Stack: Who Benefits From It', *BloombergQuint*, 20 September 2018, https://www.bloombergquint.com/opinion/niti-aayogs-national-health-stack-who-benefits-from-it (accessed 11 Mar. 2020).

NITI Aayog (2018a), 'National Strategy for Artificial Intelligence #AIForAll', June 2018, https://niti.gov.in/writereaddata/files/document_publication/NationalStrategy-for-AI-Discussion-Paper.pdf (accessed 10 Mar. 2020).

NITI Aayog (2018b), *National Health Stack: Strategy and Approach*, Delhi: NITI Aayog, https://niti.gov.in/writereaddata/files/document_publication/NHS-Strategy-and-Approach-Document-for-consultation.pdf (accessed 11 Mar. 2020).

NovoJuris Legal (2019), 'India: Cyber-Security: The Vulnerability Of Medical Institutions To Cyber-Attacks', *Mondaq*, 29 April 2019, https://www.mondaq.com/india/Technology/801952/Cyber-Security-The-Vulnerability-Of-Medical-Institutions-To-Cyber-Attacks (accessed 11 Mar. 2020).

Paul, Y., Hickok, E., Sinha, A. and Tiwari, U. (2018), *Artificial Intelligence in the Healthcare Industry in India*, The Centre for Internet and Society, https://cis-india.org/internet-governance/files/ai-and-healtchare-report (accessed 11 Mar. 2020).

Perry, P. (2016), 'How Artificial Intelligence will Revolutionize Healthcare', Big Think, 10 May 2016, https://bigthink.com/philip-perry/how-artificial-intelligence-will-revolutionize-healthcare (accessed 30 Mar. 2020).

Powell, A. C., Tyagi, H. and Ludhar, J. K. (2018), 'Digitising Indian Healthcare Records', *ISBInsight*, August 2018, https://isbinsight.isb.edu/digitising-indian-healthcare-records (accessed 11 Mar. 2020).

Press Information Bureau (2019), '"The National Digital Health Blueprint (NDHB) is just a beginning; it gives us a direction and a path to move towards strengthening healthcare delivery systems": Dr Harsh Vardhan', 29 October 2019, https://pib.gov.in/newsite/PrintRelease.aspx?relid=194063 (accessed 11 Mar. 2020).

Press Trust of India (2017), 'Healthcare IT Cos raise $5.5 bn in VC funding in January–September 2017', *Moneycontrol*, 23 October 2017, https://www.moneycontrol.com/news/business/economy/healthcare-it-cos-raise-5-5-bn-in-vc-funding-in-january-september-2017-2417471.html (accessed 11 Mar. 2020).

Purandare, V. (2018), 'Cyber attack on MGM Hospital', *Asian Age*, 19 July 2018, https://www.asianage.com/metros/mumbai/190718/cyber-attack-on-mgm-hospital.html (accessed 11 Mar. 2020).

Rajkomar, A., Dean, J. and Kohane, I. (2019), 'Machine Learning in Medicine', *New England Journal of Medicine*, 4 April 2019, 380(14): pp. 1347–58, doi:10.1056/NEJMra1814259 (accessed 10 Mar. 2020).

Rao, N. (2018), 'Who Is Paying for India's Healthcare?', *The Wire*, 14 April 2018, https://thewire.in/health/who-is-paying-for-indias-healthcare (accessed 10 Mar. 2020).

Ross, C. and Swetlitz, I. (2017), 'IBM pitched its Watson supercomputer as a revolution in cancer care. It's nowhere close', *STAT*, 5 September 2017, https://www.statnews.com/2017/09/05/watson-ibm-cancer (accessed 11 Mar. 2020).

Safi, S., Thiessen, T. and Schmailzl, K. J. G. (2018), 'Acceptance and Resistance of New Digital Technologies in Medicine: Qualitative Study', *JMIR research protocols*, 7(12), doi:10.2196/11072 (accessed 11 Mar. 2020).

SAS (undated), 'Natural Language Processing (NLP) – What is it and why it matters', https://www.sas.com/en_us/insights/analytics/what-is-natural-language-processing-nlp.html (accessed 5 Mar. 2020).

Schmelzer, R (2019), 'Understanding Explainable AI', *Forbes*, 23 July 2019, https://www.forbes.com/sites/cognitiveworld/2019/07/23/understanding-explainable-ai/#4293beac7c9e (accessed 5 Mar. 2020).

Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S. and Vertesi, J. (2019), 'Fairness and Abstraction in Sociotechnical Systems', *Proceedings of the Conference on Fairness, Accountability, and Transparency*, January 2019, pp: 59–68, doi:10.1145/3287560.3287598 (accessed 10 Mar. 2020).

Selvaraj, S., Farooqui, H. H. and Karan, A. (2018), 'Quantifying the financial burden of households' out-of-pocket payments on medicines in India: a repeated cross-sectional analysis of National Sample Survey data, 1994–2014', *BMJ Open*, 8(5), doi:10.1136/bmjopen-2017-018020 (accessed 10 Mar. 2020).

Shou, D. (2019), 'The Next Big Privacy Hurdle? Teaching AI to Forget', *WIRED*, 6 December 2019, https://www.wired.com/story/the-next-big-privacy-hurdle-teaching-ai-to-forget (accessed 10 Mar. 2020).

Siddiqui, F. R. (2008), 'Denied treatment, Dalit woman dies', *The Times of India*, 25 April 2008, https://timesofindia.indiatimes.com/city/lucknow/Denied-treatment-Dalit-woman-dies/articleshow/2980771.cms (accessed 11 Mar. 2020).

Singh, S. (2020), 'Google, Microsoft circle as India mulls extracting value from health data of 1.3 billion citizens', *The Ken*, 6 January 2020, https://the-ken.com/story/google-microsoft-india-pdp-health-data-sharing (accessed 10 Mar. 2020).

Sinha, S. (2018), 'How Microsoft Is Expanding Its Healthcare Initiative In India Using AI', *Analytics India Magazine*, 20 March 2018, https://analyticsindiamag.com/how-microsoft-is-expanding-its-healthcare-initiative-in-india-using-ai (accessed 10 Mar. 2020).

The Times of India (2019), 'Google Assistant is now available in eight more Indian languages', *The Times of India*, 25 February 2019, https://timesofindia.indiatimes.com/gadgets-news/google-assistant-is-now-available-in-four-more-indian-languages/articleshow/68152781.cms (accessed 11 Mar. 2020).

Trist, E. (1981), 'The evolution of socio-technical systems', *Occasional paper*, 2, https://stsroundtable.com/wp-content/uploads/The-Evolution-of-Socio-Technical-Systems-Trist.pdf (accessed 10 Mar. 2020).

Ubenwa (undated), https://www.ubenwa.ai (accessed 16 Mar. 2020).

USAID (2019), 'Artificial Intelligence in Global Health: Defining a Collective Path Forward', https://www.usaid.gov/cii/ai-in-global-health (accessed 5 Mar. 2020).

Vempati, S. S. (2016), *India and the Artificial Intelligence Revolution,* Research Paper, Delhi: Carnegie Endowment for International Peace, https://carnegieendowment.org/files/CP283_Vempati_final.pdf (accessed 11 Mar. 2020).

Vidyut (2018), 'French security researcher shows how to bypass password protection on mAadhaar app AGAIN', *Medianama*, 13 March 2018, https://www.medianama.com/2018/03/223-french-security-researcher-shows-how-to-bypass-password-protection-on-maadhaar-app-again (accessed 11 Mar. 2020).

Vincent, J. (2019), 'AI is worse at identifying household items from lower-income countries', *The Verge*, 11 June 2019, https://www.theverge.com/2019/6/11/18661128/ai-object-recognition-algorithms-bias-worse-household-items-lower-income-countries (accessed 16 Mar. 2020).

Wahl, B., Cossy-Gantner, A., Germann, S. and Schwalbe, N. R. (2018), Artificial intelligence (AI) and global health: how can AI contribute to health in resource-poor settings? *BMJ Global Health,* 3(4), doi:10.1136/bmjgh-2018-000798 (accessed 30 Mar. 2020).

Wickens, C. D., Clegg, B. A., Vieane, A. Z. and Sebok, A. L. (2015), 'Complacency and Automation Bias in the Use of Imperfect Automation', *Human Factors*, 57(5): pp. 728–39, doi:10.1177/0018720815581940 (accessed 10 Mar. 2020).

World Bank (2019), 'New country classifications by income level: 2019-2020', 1 July 2019, https://blogs.worldbank.org/opendata/new-country-classifications-income-level-2019-2020 (accessed 5 Mar. 2020).

World Economic Forum (2018), 'Four ways AI can make healthcare more efficient and affordable', 31 May 2018, https://www.weforum.org/agenda/2018/05/four-ways-ai-is-bringing-down-the-cost-of-healthcare (accessed 5 Mar. 2020).

World Wide Web Foundation (2017), 'Artificial Intelligence: The Road Ahead in Low- and Middle-Income Countries', https://webfoundation.org/docs/2017/07/AI_Report_WF.pdf (accessed 5 Mar. 2020).

Zandi, D., Reis, A., Vayena, E. and Goodman, K. (2019), 'New ethical challenges of digital technologies, machine learning and artificial intelligence in public health: a call for papers', *Bulletin of the World Health Organization*, 97(1), pp. 1–72, doi:10.2471/BLT.18.227686 (accessed 30 Mar. 2020).

About the Authors

Claire Muñoz Parry is assistant director of the Centre for Universal Health at Chatham House, working on a range of projects relating to global health security, universal health coverage and healthy societies. Her research work focuses on technology for health, specifically applications of artificial intelligence. Claire holds degrees in international relations and security studies from the University of Wales, and a master's in development administration and management from the University of Manchester.

Dr Urvashi Aneja is co-founder and director of Tandem Research, an interdisciplinary research collective based in India, that generates policy insights at the interface of technology, society, and sustainability. She leads Tandem's *Future of Work* and *AI and Society* initiatives. Her work examines the societal impact of algorithmic decision-making systems in India. She is an associate fellow of the Asia-Pacific Programme at Chatham House.

Acknowledgments

This paper is informed by insights gathered at the third meeting of Tandem Research's AI & Society Technology Foresight Group (TFG), held in December 2018. The TFG brings together multiple stakeholders to collectively and iteratively diagnose issues and challenges pertinent to technology and society futures in India. The meeting was supported by the Friedrich-Ebert-Stiftung, a German non-profit foundation committed to the values of democracy and social justice.

Independent thinking since 1920

Cover image: Inside view of SigTuple office in Bengaluru, on 3 May 2017. The company is helping hospitals and healthcare centres to improve the speed and accuracy of blood reports.

Photo credit: Mint/Getty Images

ISBN 978 1 78413 394 8

This publication is printed on FSC-certified paper.

designbysoapbox.com

The Royal Institute of International Affairs
Chatham House
10 St James's Square, London SW1Y 4LE
T +44 (0)20 7957 5700 F +44 (0)20 7957 5710
contact@chathamhouse.org www.chathamhouse.org

Charity Registration Number: 208223