



Operationalizing health data governance for AI innovation in low-resource government health systems: a practical implementation perspective from Zanzibar

Tracey Li¹ , Abbas Wandella¹, Richard Gomer² and Mohamed Habib Al-Mafazy³

¹D-tree, Zanzibar, Tanzania

²School of Electronics and Computer Science, University of Southampton, Southampton, UK

³Information and Communications Technology Unit, Ministry of Health, Zanzibar, Tanzania

Corresponding author: Tracey Li; Email: tli@d-tree.org

Received: 27 November 2023; **Revised:** 30 April 2024; **Accepted:** 30 April 2024

Keywords: government health systems; health data governance; low and middle income countries; policy implementation

Abstract

Improved health data governance is urgently needed due to the increasing use of digital technologies that facilitate the collection of health data and growing demand to use that data in artificial intelligence (AI) models that contribute to improving health outcomes. While most of the discussion around health data governance is focused on policy and regulation, we present a practical perspective. We focus on the context of low-resource government health systems, using first-hand experience of the Zanzibar health system as a specific case study, and examine three aspects of data governance: informed consent, data access and security, and data quality. We discuss the barriers to obtaining meaningful informed consent, highlighting the need for more research to determine how to effectively communicate about data and AI and to design effective consent processes. We then report on the process of introducing data access management and information security guidelines into the Zanzibar health system, demonstrating the gaps in capacity and resources that must be addressed during the implementation of a health data governance policy in a low-resource government system. Finally, we discuss the quality of service delivery data in low-resource health systems such as Zanzibar's, highlighting that a large quantity of data does not necessarily ensure its suitability for AI development. Poor data quality can be addressed to some extent through improved data governance, but the problem is inextricably linked to the weakness of a health system, and therefore AI-quality data cannot be obtained through technological or data governance measures alone.

Policy Significance Statement

This article discusses health data governance for artificial intelligence (AI) innovation in the context of low- and middle-income country (LMIC) government health systems, from an operational perspective. We highlight some of the practical challenges that are encountered when implementing a health data governance policy in LMIC settings that do not receive adequate attention. If left unresolved, these challenges will lead to an increase in health inequity, as the benefits of health AI will be limited to contexts where health data governance, combined with other measures, supports the development of reliable, safe, and equitable AI. Our practical implementation perspective complements the more common high-level perspective, highlighting the need to recognize and address the gaps that currently exist between policy and practice.



1. Introduction

Digital health technologies are increasingly being harnessed to strengthen health systems in countries all over the world, in line with the World Health Organization's Global Strategy on Digital Health (WHO, 2021). This has led to an explosion in the amount of data that is generated within health systems (Dinov, 2016). Governments and researchers are eager to use that data to conduct analyses and research and to create artificial intelligence (AI) models that can radically transform healthcare, leading to improved health outcomes. However, it is becoming apparent that numerous challenges must be overcome in order for health data to be used to its maximum potential. Technological barriers are rapidly being lowered, thanks to a proliferation of off-the-shelf software tools and on-demand computing infrastructure (Harvard Business Review, 2023), while issues related to bias and fairness in the development and deployment of AI are currently being widely discussed (Agarwal et al., 2022). Alongside these are questions of how to evaluate, regulate, and increase public understanding of health AI (Kelly et al., 2019; Singh et al., 2020). It is now widely acknowledged that appropriate governance is necessary if we are to create AI that is reliable, safe, and trustworthy (Shneiderman, 2020), as well as fair and equitable (d'Elia et al., 2022), rather than AI which is optimized solely for accuracy and efficiency. Both the governance of the development and use of AI, and the governance of the data that is used to develop AI models, are necessary. It is the second topic—health data governance for AI innovation—that is the focus of this paper.

Data governance is defined by DAMA International as the “exercise of authority and control over the management of data” (DAMA International, 2009). Some sources differentiate between governance and management, using “governance” to refer to what decisions must be made to ensure effective management and use of information technology (and by extension data) and who makes those decisions, whereas “management” involves making and implementing decisions (Weill and Ross, 2004). Many non-academic practitioners use the term “governance” to refer jointly to what decisions must be made and by whom, as well as the implementation of those decisions, and think of data governance as the practices that ensure the security, privacy, accuracy, availability, and usability of data (Google Cloud, 2023). Throughout this work, we have chosen to adopt this latter approach and use the term “data governance” to encompass both decision-making and implementation.

In the global health context, there is a strong emphasis on the ethical component of data governance and global advocacy efforts focused on driving improvements (Transform Health Coalition, 2023a; WHO, 2023). These global efforts typically result in high-level principles which can be generally applied in all contexts (OECD, 2017; PAHO, 2021; Transform Health Coalition, 2023b). However, the practical implementation of those principles can look very different, and be accompanied by distinct challenges, in different contexts. For example, a setting in which digital health programs exist but where technology is otherwise scarce presents challenges that are different to a setting in which digital infrastructure and technology are ubiquitous. Therefore, although conversations about health data governance commonly focus on policy and regulation, there is a need to progress to discussions about the practical, and even more importantly pragmatic, implementation of those policies and regulations. Effective, localized implementations are necessary to produce the meaningful real-world changes that we are aiming for—namely improvements to people's privacy and health outcomes—rather than changes that exist only as theoretical ideals written on paper. Yet the implementation aspect receives minimal attention, although it is fraught with enormous challenges. The European Union's experience of implementing the General Data Protection Regulation (GDPR) has demonstrated that for both the regulators tasked with enforcing the GDPR, and the businesses and organizations that have to comply with the regulation, implementation is extremely complex and costly (Sirur et al., 2018). Although the GDPR is much broader than a health data governance regulation and therefore necessarily more complicated, the experience still demonstrates the difficulty of large-scale implementation.

In order to meaningfully focus on the practical aspects of health data governance, we will restrict our discussion to three specific topics—informed consent, data access and security, and data quality—rather than attempting to comprehensively address all of data governance at a comprehensive level. Similarly, we will not attempt to address all types and uses of AI, but instead restrict our study to machine learning algorithms that are trained on patient data to provide diagnostic or clinical outputs that support clinical

decision-making (Adlung et al., 2021; Awaysheh et al., 2019; Panch et al., 2018; Peiffer-Smadja et al., 2020). This excludes a broad range of other technologies, such as large language models that can be used to create chatbots that provide health information or support research (Sallam, 2023) and computer vision for medical imaging (Esteva et al., 2021). Since these technologies are distinct in terms of both the underlying AI models and the data required for training those models, we believe each requires a separate discussion. Therefore, we choose to focus solely on the first one in this paper. Also, we will focus on the data generated within government health systems, as opposed to data produced in health programs that are run by private or third sector actors, although much of our discussion is applicable to all cases.

The content of this paper is primarily informed by work conducted in Zanzibar by D-tree, a non-government organization. D-tree has been supporting the Zanzibar Ministry of Health (ZMOH) for over 10 years to strengthen the Zanzibar health system through the use of technology and data and, since 2022, has dedicated specific effort to support ZMOH to improve health data governance at an operational level, with some advisory support from academic researchers. The paper is based on the primary authors' first-hand experience of working within the Zanzibar health system, where we have either observed and explored, or started to directly address, the issues that we discuss in this work. We believe that many of our learnings and insights are applicable to low- and middle-income country (LMIC) government health systems in general and use our case study in Zanzibar to concretely illustrate many salient points. We also believe that although this paper is framed around health data governance for AI innovation, much of our discussion is applicable to research and evaluation more broadly, not just to AI.

The following section is divided into three parts. In the first part, we discuss the topic of informed consent at both a conceptual and practical level, with reference to some specific details in Zanzibar. In the second part, we describe our experience of making operational improvements to data access management and data security within the Zanzibar health system. In the third part, we discuss the data quality challenges that are present in Zanzibar and many other LMIC settings. We close the paper with a section summarizing our key learnings.

2. From policy to practice

2.1. Informed consent

Informed consent is one of the core principles of many data governance frameworks. However, the practice of obtaining informed consent is frequently both literally and metaphorically reduced to a box-ticking exercise, due to the difficulties of communicating about data use (Custers et al., 2018), designing an interaction in which meaningful consent can be sought, and, particularly relevant to LMIC settings, documenting consent in cases where it is not possible to obtain written consent.

In many cases, obtaining consent to administer medical treatment and obtaining consent to use data are distinct and come with different legal requirements, yet the requests often happen at the same time. It is easy for requests to those multiple consents to become conflated, or for the distinction between them to be overlooked. Somebody may wish to consent, urgently, to medical treatment, while being too distracted by stress or pain to attend meaningfully to a decision about use of their health data. Thus, it is necessary to clearly distinguish between medical consent and data consent. Consent to data processing—particularly secondary data processing like AI applications—is perhaps the most easily overlooked, given that it is the most distant and unfamiliar action that consent is typically sought for, but it is also the type of consent that we are most concerned with in this paper.

There are two key motivations to attend to consent—the first is legal and the second is ethical. Legally, consent plays a key role in many data protection regimes, most notably the European Union's GDPR, and even in jurisdictions without such legislations—such as Zanzibar—international data flows and processing mean that some datasets come into the scope of those laws later in their lifecycle. Ethically, consent is “morally transformative” (Hurd, 2018). It can remove moral barriers to specific conduct, “making permissible what was otherwise prohibited; making right what was otherwise wrong,” and the idea that data about people should be used only with their consent is a fairly widespread normative belief,

particularly for secondary purposes such as healthcare research where the purpose of processing may not be the purpose for which the data was originally collected, and hence may not be a purpose that the data subject was expecting. Such ethical concerns are a key reason to engage with consent even when operating in circumstances where formal legal obligations do not require it. We acknowledge that several data protection legislations include exemptions to reuse data for secondary purposes, without explicit consent, when the reuse will be for public benefit. However, we will not consider this here as defining what constitutes “public benefit” requires a broad discussion that extends beyond the scope of our current work.

Obtaining meaningful consent is challenging for multiple reasons. Consent is an inherently complex act based on two parties having mutual and meaningful understanding and agreement. Therefore, assessing whether consent has arisen during an interaction, or designing an interaction to reliably result in consent, is highly non-trivial. A general-purpose set of principles for consent-building (Friedman et al., 2002) comprises “disclosure” (information provided should be accurate and attempt to disclose foreseeable harms), “comprehension” (an individual needs to accurately interpret what is being disclosed), “voluntariness” (an individual’s action must not be controlled or coerced), “competence” (the individual must have the capability required to give consent), and “agreement” (there must be a clear opportunity to accept or decline what is proposed). Implementing each principle comes with a number of challenges. For example, deference to perceived authority, or even just politeness, can undermine the voluntariness of consent. Relatively subtle factors in an interaction—such as whether a healthcare worker seems to present a request for consent as a serious decision or as a mere formality—can influence the degree of engagement that a patient gives to the decision, and even whether they feel compelled to agree. Therefore, seemingly small differences in the implementation of the consent process can have a large impact on people’s decisions (Utz et al., 2019). With respect to competence, one group (Manson and O’Neill, 2007) notes that “incompetence and impaired competence to consent are more common in medical practice than elsewhere, since impaired cognitive capacities are a common effect of illness and injury.” In terms of disclosure and comprehension, there is a clear need for more research into how to communicate with patients about their data (NHS Confederation, 2024). Communication can be especially difficult in LMICs because the level of digital literacy is low (World Bank, 2021), meaning that many people are unfamiliar with concepts and terms that are now commonly used in high-income settings. Additionally, some people may lack the literacy skills required to engage with written materials, meaning that all information has to be provided verbally and that consent needs to be obtained verbally rather than in written form, which requires more effort to be invested in the documentation of that consent by, for example, audiovisual means (Benitez et al., 2002). Also, much of the vocabulary associated with data use and data governance is not easily translatable into non-English languages (Datasphere Initiative, 2023) and studies have shown that it is necessary for communication to be closely tailored to each specific context (Cheah et al., 2018) meaning that a “one size fits all” approach is not viable.

Despite the challenges, it is necessary to make progress with improving how informed consent is sought if we are to develop reliable, safe, and equitable AI for everyone. In LMIC settings, the challenge of doing so is the greatest, but the need for better data governance is urgent. For example, a rural village in sub-Saharan Africa, such as villages in Zanzibar, may not have reliable electricity or internet connection, but can be served by community health workers who are equipped with smartphone-based digital tools that support them to deliver critical health services and to record vital health information (Owoyemi et al., 2022). That information includes details about each patient’s health and may also include data that directly or indirectly pertains to a patient’s race, ethnicity, religious beliefs, or sexual orientation. Under the remit of many data protection regulations, these data are classified as highly sensitive and are subject to the strictest standards of protection (ICO, 2024). Doubly so because many of the recipients of essential health services belong to “vulnerable” groups—the term used to describe groups that are at higher risk of harm than the general population, such as young children, pregnant women, and people living with HIV. However, many health workers in LMICs do not receive training about the collection and management of data, despite being obliged to conduct those tasks alongside their primary function of delivering essential healthcare services, all with very limited time and resources (Ngusie et al., 2021; Nwankwo and Sambo, 2018; Siyam et al., 2021). As a result, there is limited awareness about why data is collected from patients, how it will be

used, who will have access to it, and how it should be managed. Yet it is necessary for both patients and health service providers to understand exactly these points if we expect to be able to obtain informed consent.

As well as a lack of understanding about the proper management of data, the need to obtain consent for data use is also often not understood, neither by patients nor data collectors. Even when it is understood, the parties responsible for collecting the data may choose not to ask for consent due to a number of reasons. These include an assumption that everyone will anyway provide consent because the benefits “obviously” outweigh the risks, an assumption that as long as data is de-identified before being shared there will not be any concerns (McKay et al., 2023), and concerns that people will become unnecessarily worried about the risks and subsequently withdraw from the program or study. To test these assumptions, D-tree conducted a study in Zanzibar (unpublished), administering a survey to 97 people enrolled in the community health program to gauge their understanding and attitudes regarding the data that is collected about them as part of the program. Our key findings were that 44 percent of respondents were not aware of who could access their data and 32 percent did not know why their data was collected. We did not follow up with respondents who answered affirmatively to either or both questions to check that they had a correct understanding, and therefore it is possible that the percentages of negative responses are actually higher. We found that respondents were very accepting of their data to be shared with health program staff, health facility staff, and government stakeholders, with only one percent of respondents expressing concern in this respect. The percentage of respondents with concerns increased when asked whether they would be happy for their de-identified data to be shared with university researchers either within the country (four percent expressed concerns) or internationally (five percent). Respondents expressed concerns related to confidentiality, the large geographical distance that would be between themselves and their data, shame about their health conditions meaning that they did not want others to know about it, and not understanding how sharing data would benefit them. Over 95 percent of respondents were receptive to receiving more information about how their data was being used by various stakeholders. We therefore conclude that although a large proportion of respondents lack basic understanding about who has access to their data and why, they are generally comfortable for their data to be shared with local relevant stakeholders, with a few concerns arising when data is to be shared with researchers. We therefore believe that we can improve communication, in such a way that people who are enrolled in the program can better understand what is happening to their data and provide meaningful consent, without causing undue fears that lead to disengagement from the program.

Although our discussion has so far focused on individual consent, it is worth considering wider public perspectives and in particular the public acceptability of health data use. A social license is typically required before public health data can be acceptably repurposed (Paprica et al., 2019) and the absence of such a license has been implicated in the high-profile failure of schemes such as the UK’s “care.data” (Sterckx et al., 2016). Consent, as typically understood as an individual decision, may in itself be one means of improving public acceptability, but other authors point to the conceptual validity of consent as a collective action (Varelius, 2008) inherently linked to constructs such as social license and participatory or democratic decision-making. For the most part, we have ourselves considered consent as an individual act because that is how it exists legally in most jurisdictions, but a collective view may have some benefits in LMIC contexts, for example, as a means to work with populations that have lower individual data literacy, either instead of or in addition to individual consent. Collective approaches to consent are consistent with the concept of privacy as a public good (Fairfield and Engel, 2017) and contiguous with collective data governance mechanisms such as those that might be deployed within data institutions or data co-ops (Gomer and Simperl, 2020).

As a final note, we will highlight that obtaining consent to use data for AI applications necessitates not only communicating about general concepts related to data use but also communication about the AI technology itself. We could not find any studies that have been conducted about perceptions and understanding of AI in a LMIC context, so will refer instead to a nationally representative survey in the UK that explored people’s views about the use of AI within the National Health Service (NHS) (NHS England Transformation Directorate, 2023). Around 40 percent of respondents wanted “Greater

transparency about how AI works,” “Clear and accessible information on how data about me is used,” and “Clear and accessible information about which private companies are working with the NHS and social care to develop AI technologies.” This demonstrates that even in a setting like the UK where the use of AI is already ubiquitous, there is still a long way to go in terms of increasing public understanding to the level where people can make truly informed decisions about how their data is used for AI development. Meaningful consent for AI use is predicated, to some degree, on background knowledge that many people do not yet have. In a LMIC setting where a large proportion of the population may not have exposure to even relatively basic technology such as smartphones, are not familiar with what AI is or how it works, and therefore have almost certainly not considered the risks posed by having their health data stored in a research database, we must acknowledge that a huge amount of progress is required before informed consent for health AI can be meaningfully implemented.

2.2. *Data access and data security*

Governments are the owners of some of the biggest and potentially most valuable health datasets in the world. Yet while there is substantial investment in developing the technology to collect that data, there is minimal investment in building capacity within governments to manage the data in a manner that facilitates effective usage, including AI innovation, while respecting all legal and ethical requirements. When efforts are made, they often start and end with the creation of policies but go no further. Subsequently, there is limited understanding about what to do with a policy once it has been written, specifically how to translate it into real-life actions that result in the improvements to individuals’ privacy and health that we are seeking. In this section, we report on the process that was undertaken within the Zanzibar health system to address this gap, through an ongoing collaboration between the ZMOH and D-tree, an international NGO that has been providing technical support to ZMOH for many years.

In 2022, ZMOH commissioned a consultant to create the ministry’s first Data Protection and Sharing policy as part of Zanzibar’s Digital Health Strategy (Revolutionary Government of Zanzibar, 2020). The lack of follow-on actions after creating the policy was complicated by limited ownership, as responsibilities for data management are divided between the Information and Communications Technology (ICT) and Health Management Information Systems (HMIS) units. In order to address this, ZMOH shared the policy with D-tree who then convened stakeholders from both the ICT and HMIS units, plus other relevant units, to build an understanding of why it is necessary to operationalize the Data Protection and Sharing policy and to agree on how to work together on this task. Once consensus had been reached, a ZMOH committee was created to champion the operationalization process. D-tree led a series of capacity-building workshops for the committee members to develop their understanding of core concepts related to data governance, the content of the policy, and to create a plan to operationalize the policy. Given the limited resources available, it was essential for the committee to decide which parts of the policy to prioritize since operationalizing the entire policy in the immediate future would not be feasible. The priorities which emerged were improving information security and the management of data access. The committee reported that an absence of formal processes resulted in ad hoc decisions about who can access data and concerns about the security of the data, particularly in the case of data access requests from external parties. Additionally, there were situations in which staff who needed to access data to fulfil their work duties were unable to do so. Concerns about lack of accountability were frequently raised.

Once funding had been mobilized, D-tree initiated two parallel work streams—one focused on data access management and one on information security. For the first, D-tree drafted a set of data access management guidelines. The guidelines included descriptions of the roles and responsibilities that need to be assigned in order to effectively manage data access, with a template included to assign those roles for each dataset. Another template facilitated the classification of data types in terms of their confidentiality level. The main body of the guidelines then comprised detailed procedures for requesting, approving, reviewing, and maintaining data access for ZMOH staff, and separately for external data users. A data request form template and data use agreement template for external users were included as an appendix,

with the agreement including clauses to ensure that the data is only used for the purpose that ZMOH has agreed to and by the individuals that ZMOH has approved, and that no part of the data or any information derived from the it can be disclosed without prior approval from ZMOH. The final appendix provided technical guidance for technical staff at ZMOH to de-identify data using techniques such as pseudonymization and perturbation, prior to data being shared. D-tree circulated the initial draft among stakeholders from the ICT, HMIS, Health Information Systems, and Monitoring and Evaluation units prior to a workshop. D-tree collected feedback from that workshop and incorporated it into a revised draft, which D-tree and ZMOH then circulated amongst a wider group of stakeholders, including technical officers from all programmatic units (Zanzibar Integrated HIV, TB and Leprosy Program, Non-Communicable Diseases, Neglected Tropical Diseases, Zanzibar Malaria Elimination Program, Integrated Reproductive and Child Health, the Health Promotion Unit, the Chief Pharmacist Office, and Mnazi Mmoja—Zanzibar’s referral hospital) and relevant government bodies (the e-Government Agency, the Office of Chief Government Statistician—the national statistics office, and legal advisors). D-tree incorporated the feedback from that wider circle of stakeholders into an updated draft which was presented to the Zanzibar Health Sector Performance Technical Working Group for their review and endorsement. The final version was then translated by a technical translator into Swahili. Both the Swahili and English versions of the document were circulated by D-tree for a final review before being presented to the ZMOH senior leadership team for their endorsement and approval. The senior leadership team of ZMOH were kept informed at all stages of the process to ensure the necessary high-level buy-in and engagement. At the time of writing (March 2024), the guidelines have been printed and are being disseminated, and D-tree plans to conduct a series of training sessions for ZMOH staff and other government stakeholders to facilitate understanding and compliance with the guidelines.

Regarding the information security work stream, D-tree supported ZMOH to hire a consultant to conduct a situation analysis of ZMOH’s systems and then to create comprehensive information security guidelines that can be generally applied across all systems. Once the consultant had been contracted, the ZMOH senior leadership team met with the consultant to initiate the assignment. The consultant conducted an assessment of key ZMOH systems, including interviews with key stakeholders. They identified and documented issues and recommendations, which were presented to ZMOH. The consultant also gathered and reviewed existing ICT standards and guidelines and updated them to improve their effectiveness and relevance. Alongside this, the consultant developed a training manual to support staff to comply with the updated guidelines. The new standards and guidelines were officially approved by ZMOH leadership, enabling their formal adoption. Finally, the consultant conducted a week-long training-of-trainers, equipping a team of ZMOH staff to train their peers on the guidelines.

The entire process described above took approximately 18 months. After the completion of training for both sets of guidelines, we expect to see evidence of increased efficiency, transparency, and accountability during data sharing because staff will be following a defined process that includes documenting and signing off on critical decisions. We also expect to see systematic adherence to defined security standards. These improvements will contribute to increased privacy and accountability to the patients of Zanzibar’s health system.

Our implementation has currently focused predominantly around processes and people, with limited focus on technology. This approach is driven by the current needs and constraints of the health system. The needs are to improve the flow and security of data within the health system to facilitate operational efficiency while respecting patient confidentiality and to be able to respond to requests from outside the health system to access data for the purposes of analysis and research—which necessitates being able to judge which requests should be granted and then providing access to data in an appropriate manner while respecting patient confidentiality. The first need can be addressed to a large degree by the introduction of processes to standardize how data access is provisioned to health system staff, alongside assigning specific individuals to be responsible for deciding what data each staff member needs to access, reviewing that access periodically, and implementing the technical measures necessary to provision access (typically creating a user account within existing systems and appropriately configuring permissions). This is what is included in the newly produced guidelines and forthcoming training. For the second need, the decision

about which data access requests to approve is currently undertaken by ZMOH staff, who may consult with relevant colleagues or partners to be able to reach a decision, and the new guidelines should ensure that this is done in a transparent and systematic manner. There is room in the future to explore the creation of a formal data access committee (DAC) to review data access requests that includes representatives from the community, from where the data originates, as well as stakeholders from within the health system. Such DACs are becoming increasingly common in high-income contexts, although there is not yet any generally accepted framework for their organization (Cheah and Piasecki, 2020). It will therefore be pertinent to investigate what would be the optimal composition and mode of operation for a DAC in Zanzibar and the factors contributing to that. Regarding the need to provision data access to external researchers, there are various technological solutions to address this, such as data enclaves (Lane and Schur, 2010) and federated learning systems (Rieke et al., 2020). These technologically advanced solutions afford many benefits in terms of automation and enhanced privacy, but the high technical complexity necessitates significant resources and ongoing investment to implement and then maintain those solutions. As sufficient resources are not currently available in Zanzibar, it is therefore necessary to instead consider less complex and lower cost approaches at the immediate time, with a view to iteratively improving as more resources become available so that technologically complex solutions can realistically be introduced and maintained.

We will note also that our implementation has not directly addressed the question of how to ensure that the population and health system of Zanzibar will benefit from the use of data by external researchers, especially those who are based outside the country (from where the vast majority of requests originate). From a pure research perspective, involving local research collaborators and maintaining equity in the collaboration is essential (Hedt-Gauthier et al., 2018). However, there can also be an expectation for research to result in actionable insights that will benefit the country, but with researchers and non-researchers having a different understanding of what this constitutes. Additionally, there are many challenges related to translating research into practice (Glasgow and Emmons, 2007) which means that many research studies typically end as academic publications, that cannot be easily understood by non-academic readers, that contribute to the scientific knowledge base but do not produce any direct tangible benefit to the health system. It is therefore necessary for ZMOH to define what benefit they wish Zanzibar to receive and factor this into decisions about which data access requests to grant, and specify terms that are most likely to ensure the manifestation of that benefit. To this end, ZMOH has previously requested academic institutions to formally sign a collaboration agreement, which extends to a broader scope than the data use agreement, that obligates the researcher to contribute to building research capacity within ZMOH in order to guarantee some tangible benefit.

We believe that the challenges described here are not unique to Zanzibar, but are instead common in many LMIC settings. Our experience demonstrates the need to increase awareness, capacity, and resources within governments to enable the operationalization of data governance policies. Extensive sensitization and training are required to build an understanding of the importance of data governance and the necessity to plan follow-up actions after the creation of a policy in order to operationalize it. The operationalization process necessitates engagement with a diverse set of stakeholders in order to ensure effective and contextually appropriate implementation, as well as buy-in from key stakeholders. The limitation of resources in LMIC government systems leads to a requirement to be pragmatic and to address the priorities that will have the biggest impact, rather than taking an idealistic approach of trying to address all issues and subsequently not substantially addressing any. Finally, we note that the specialist technical expertise needed to support operationalization is, in general, not readily available to LMIC governments. People with the required technical skills are unlikely to work within a government setting because they can obtain a significantly higher salary and more favourable working conditions elsewhere, and technical consultants are expensive and can be difficult to source. Additionally, there are very few materials available to provide guidance for this work that are relevant, accessible, and actionable to staff in a government health system. To at least partly address this, we believe that peer-to-peer learning exchanges between governments can be an effective way to accelerate progress, whereby governments who have made progress in certain areas can exchange their learnings with peers who are working in a similar

context. D-tree, together with the Global Partnership for Sustainable Development Data, facilitated one such learning exchange between the ministries of health of Zanzibar and Kenya which included discussions about challenges, successes, and potential areas for cross-government collaboration (Global Partnership for Sustainable Development Data, 2023). Such exchanges should also serve to reduce the degree to which governments have to rely on support from development partners, affording greater independence and autonomy to governments.

2.3. Data quality

Health systems are complex and generate huge amounts of data, especially when digitally enabled. This includes data about finances, human resources, logistics, supplies and equipment, and service delivery. It is pertinent to consider how to reuse that data for AI development, research, and analysis, given the availability of tools and technologies to extract insights from the data. Therefore, many data governance frameworks promote the reuse of data for innovation (OECD, 2017; Transform Health Coalition, 2023).

As mentioned in the introduction, we restrict our attention to patient-level service delivery data, which can be used to create AI models that output diagnoses or predictions (Panch et al., 2018). There is a lot of excitement about the potential to use patient's health data to predict their future health outcomes, thereby enabling more personalized and timely care. Implicit in this assumption is that the data will have sufficient predictive power to be able to make a prediction at a level of accuracy that is clinically useful, that it will do so fairly across all population groups, and that the prediction can be acted upon in a manner that will improve the client's outcome. To discuss these assumptions, it will be instructive for us to first consider the case of data that is collected specifically for the purpose of AI development, and we will then contrast that with the case of data that is routinely collected as part of service delivery in Zanzibar.

When data is purposefully collected for the development of health AI, the collection can be conducted in such a way to maximize the suitability of the dataset for its intended purpose (Lacuna Fund, 2023). The collection can be designed to capture as much data with as much relevant detail as is practically possible. Also, the representativeness of the data is of high priority from the perspective of equity, and there are now standardized guidelines (STANDING Together, 2023), which extend upon more general guidelines for documenting machine learning datasets (Gebu et al., 2021), to ensure that biases in health AI datasets are minimized, or that the biases are at least documented appropriately (although it should be noted that the understanding of terms such as "gender," "race," and "ethnicity" varies across cultures and therefore limits the degree to which standardization is possible). We therefore see that there are many advantages to collecting data specifically for AI development. On the other hand, the disadvantage is that the collection is conducted as a one-off event, and there is a risk that "data drift" can occur over time. This is a phenomenon where the relation between the input data and the AI model outputs changes over time, due to a number of factors that cannot be precisely predicted or quantified. This can make AI model outputs increasingly inaccurate as time progresses (Sahiner et al., 2023). The only way to mitigate this is to repeat data collection periodically to ensure that data is up-to-date, and that AI models are periodically re-trained on the most recent data. This repeated data collection exercise requires ongoing investment to sustain.

Comparing the above case with data that is being continuously collected as part of ongoing health programs, in the case of routinely collected data we remove the obstacle of data drift, as well as the costs associated with data collection. However, as the data is collected for the purpose of enabling effective health service delivery, not specifically for AI development, there will inevitably be some shortcomings. In terms of the representativeness of the data, only the people reached by the health program will appear in the dataset, and this reach is often determined by the limited resources that are available. This impacts health equity (d'Elia et al., 2022). Therefore, it cannot automatically be assumed that just because data exists, and that it is a large dataset covering a large proportion of a country's population, it is suitable for AI development. For this reason, it is necessary to ensure that the dataset is comprehensively and transparently documented, particularly highlighting areas of known or potential bias that are known to be, or that might be relevant to healthcare. This enables the user of the data to judge what they can and cannot use the data for, the extent to which AI model outputs should or should not be generalized, and what conclusions can or

cannot be drawn from the data (Gebru et al., 2021; STANDING Together, 2023). Besides the representativeness of a dataset, the predictive power of a dataset also needs to be considered. This is related to what data is collected as well as the data volume. Data collection in LMICs is often reduced to a minimum level to minimize the burden on health service staff (Siyam et al., 2021), and therefore the data will likely not be as detailed as would be ideal for AI development. Additionally, an absence of, or problems with equipment such as blood pressure monitors that are needed to generate data, alongside limited infrastructure which can make data collection arduous or at times impossible, reduce the quality of the dataset. A lack of standardization around what data to collect and report can result in inconsistent reporting, making it difficult or impossible to combine data reported from different locations. Finally, misclassification of, for example, cause of death, occurs if staff have not received sufficient training to be able to accurately identify the cause of death (Ahmed et al., 2023). Misreporting may also happen intentionally if staff perceive there will be negative consequences for them if they report truthfully, for example, in cases where a patient experiences a preventable negative outcome while under their care. All of these factors are present in Zanzibar, are widely reported in many LMICs, and affect the performance of AI models that can be produced from that data (Agency Fund, 2023), with one systematic scoping study (Ciecierski-Holmes et al., 2022), highlighting limited data availability to be one of the barriers to the development and adoption of health AI models in LMICs. Furthermore, unrelated to the data but relevant to whether an AI model should be deployed or not is the question of which AI outputs are actionable in a certain context. Specialist medical equipment or supplies are the only effective intervention for some health conditions, but those equipment and supplies are not reliably available in many LMICs. Therefore, an AI model that predicts neonatal mortality may perform at a clinically acceptable level, and also the necessary intensive care equipment to change that outcome may be available in a high income context, but it is unlikely that both of those conditions will be met in a LMIC context.

We conclude that although improvements can be made to the quality of service delivery data by providing better training to the healthcare workers who collect the data (Nwankwo and Sambo, 2018), and that providing standardized documentation to data users (STANDING Together, 2023) should improve effective usage of the data, those efforts alone may not be sufficient to bring data to a level where it is optimal for AI innovation. Reaching that level requires having a sufficient number of trained and supported healthcare workers who can provide patients with high-quality care and also collect high-quality data, ensuring that there are sufficient medical equipment and supplies so that the necessary patient data can be collected, and improving infrastructure so that all equipment that is required for data collection is fully functional. Therefore, the poor quality of data in Zanzibar, and in LMICs generally, cannot be solved by data governance alone, as the problem is inextricably linked to the weaknesses of the health system.

3. Summary

A summary of the challenges discussed in this paper is provided below. These are all based on our first-hand experience of implementing health data governance policies and principles within the Zanzibar government health system and believe that they are broadly applicable to LMICs in general. We hope that this work serves to highlight the gap between policy and practice in LMICs and encourages action to address that gap.

- Informed consent is frequently included as a guiding principle in health data governance frameworks, but the challenge of conveying meaningful understanding about data and AI, and designing an interaction in which consent can be meaningfully obtained, impedes our ability to obtain informed consent. This is especially true in LMICs where levels of exposure to technology are low and there is little familiarity with concepts and terminology related to data and AI. More attention needs to be focused on developing effective modes of communication and interaction that enable meaningful informed consent to be obtained.
- Operationalizing a data governance policy in a government health system, through the development of guidelines and standards, followed by training staff to adhere to those guidelines and standards, is

a lengthy process requiring a substantial amount of time, resources, and technical expertise that are very limited in LMICs. Extensive sensitization is required to generate awareness and understanding, and it is necessary to engage with a large and diverse set of stakeholders. Support and resources are generally lacking for this operationalization process.

- Reusing data collected in health programs for AI innovation is promoted by many governance frameworks, without taking into consideration the quality of the data. Poor data quality in LMICs is linked to weaknesses in the health system, in terms of under-staffing, lack of equipment and supplies, and unreliable infrastructure. It is therefore essential to produce comprehensive documentation for datasets to ensure that biases and other limitations are taken into account during the usage of that data. Improvements to data quality can be made by providing staff with better training and support, but obtaining AI-quality data necessitates strengthening the health system as a whole rather than approaching the problem from a purely technological or governance perspective.

Acknowledgements. D-tree is grateful to the Zanzibar Ministry of Health for its long-term collaboration and for approving the publication of this paper.

Author contribution. Conceptualization: TL. Investigation: AW (section 2.2); TL (section 2.1, 2.3). Methodology: AW (section 2.2), TL (section 2.1, 2.3). Writing original draft: TL (all sections); RG (section 2.1); AW (section 2.2). All authors approved the final submitted draft.

Provenance. This article is part of the Data for Policy 2024 Proceedings and was accepted in Data & Policy on the strength of the Conference's review process.

Funding statement. This work was supported by funding from the Patrick J. McGovern Foundation to D-tree for a program of work "to support the government of Zanzibar in maximizing the value of its health data." The funder had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interest. The authors declare no competing interests exist.

References

- Adlung L, Cohen Y, Mor U and Elinav E (2021) Machine learning in clinical decision making. *Med* 2(6): 642–665. <https://doi.org/10.1016/J.MEDJ.2021.04.006>.
- Agarwal R, Bjarnadottir M, Rhue L, ... Gao G (2022) Addressing algorithmic bias and the perpetuation of health inequities: an AI bias aware framework. *Health Policy and Technology* 12(1): 100702. <https://doi.org/10.1016/j.hlpt.2022.100702>.
- Agency Fund (2023) Early Lessons From a Precision Health Partnership. Available at <https://www.agency.fund/post/early-lessons-from-a-precision-health-partnership> (accessed 22 November 2023).
- Ahmed SM, Cresswell JA and Say L (2023). Incompleteness and misclassification of maternal death recording: a systematic review and meta-analysis. *BMC Pregnancy and Childbirth*, 23(1): 794. <https://doi.org/10.1186/s12884-023-06077-4>.
- Awaysheh A, Wilcke JR, Elvinger F, Rees LP, Fan W and Zimmerman KL (2019) Review of medical decision support and machine-learning methods. *Veterinary Pathology* 56(4): 512–525. <https://doi.org/10.1177/0300985819829524>.
- Benitez O, Devaux D and Dausset J (2002) Audiovisual documentation of oral consent: a new method of informed consent for illiterate populations. *The Lancet* 359(9315), 1406–1407. [https://doi.org/10.1016/S0140-6736\(02\)08361-7](https://doi.org/10.1016/S0140-6736(02)08361-7).
- Cheah PY, Jatupornpimol N, Hanboonkunapakarn B, ... Bull S (2018) Challenges arising when seeking broad consent for health research data sharing: a qualitative study of perspectives in Thailand. *BMC Medical Ethics* 19(1): 1–11. <https://doi.org/10.1186/S12910-018-0326-X>.
- Cheah PY and Piasecki J (2020) Data access committees. *BMC Medical Ethics* 21(1): 1–8. <https://doi.org/10.1186/S12910-020-0453-Z>.
- Ciecierski-Holmes T, Singh R, Axt M and Barteit S (2022) Artificial intelligence for strengthening healthcare systems in low- and middle-income countries: a systematic scoping review. *npj Digital Medicine* 5, 162. <https://doi.org/10.1038/s41746-022-00700-y>.
- Custers B, Dechesne F, Pieters W, Schermer B and Van Der Hof S (2018) Consent and privacy. In *The Routledge Handbook of the Ethics of Consent*. Routledge (London), 247–258.
- DAMA International (2009) *The DAMA Guide to the Data Management Body of Knowledge*. DAMA-DMBOK. Technics Publications (Sedona).
- Datasphere Initiative (2023) The Data Language. Available at <https://www.thedatasphere.org/news/the-data-language/> (accessed 20 November 2023).
- d'Elia A, Gabbay M, Rodgers S, ... Frith L (2022) Artificial intelligence and health inequities in primary care: a systematic scoping review and framework. *Family Medicine and Community Health* 10(Suppl 1). <https://doi.org/10.1136/fmch-2022-001670>.

- Dinov ID** (2016) Volume and value of big healthcare data. *Journal of Medical Statistics and Informatics* 4(1). <https://doi.org/10.7243/2053-7662-4-3>.
- Esteva A, Chou K, Yeung S, ... Socher R** (2021) Deep learning-enabled medical computer vision. *npj Digital Medicine* 4(1): 5. <https://doi.org/10.1038/S41746-020-00376-2>.
- Fairfield J and Engel C** (2017) Privacy as a public good. In Miller RA (ed.), *Privacy and Power: A Transatlantic Dialogue in the Shadow of the NSA-Affair*. Cambridge: Cambridge University Press, pp. 95–128. <https://doi.org/10.1017/CBO9781316658888.004>.
- Friedman B, Howe DC and Felten E** (2002) Informed consent in the Mozilla browser: implementing value-sensitive design. In *Proceedings of the 35th Annual Hawaii International Conference on System Sciences*. IEEE (New York), p. 10.
- Gebru T, Morgenstern J, Vecchione B, ... Crawford K** (2021) Datasheets for datasets. *Communications of the ACM* 64(12): 86–92. <https://doi.org/10.1145/3458723>.
- Glasgow RE and Emmons KM** (2007) How can we increase translation of research into practice? Types of evidence needed. *Annual Review of Public Health* 28(1): 413–433. <https://doi.org/10.1146/ANNUREV.PUBLHEALTH.28.021406.144145>.
- Global Partnership for Sustainable Development Data** (2023) Advancing Expertise in Health Data Governance: Lessons from the Ministries of Health in Kenya and Zanzibar. Available at <https://www.data4sdgs.org/resources/advancing-expertise-health-data-governance-lessons-ministries-health-kenya-and-zanzibar> (accessed 21 November 2023).
- Gomer RC and Simperl E** (2020) Trusts, co-ops, and crowd workers: could we include crowd data workers as stakeholders in data trust design? *Data & Policy* 2, e20. <https://doi.org/10.1017/dap.2020.21>.
- Google Cloud** (2023) What is Data Governance?. Available at <https://cloud.google.com/learn/what-is-data-governance> (accessed 20 November 2023).
- Harvard Business Review** (2023) How the Cloud Is Changing Data Science. Available at <https://hbr.org/2023/11/how-the-cloud-is-changing-data-science> (accessed 8 March 2024).
- Hedt-Gauthier B, Airhihenbuwa CO, Bawah AA, ... Volmink J** (2018) Academic promotion policies and equity in global health collaborations. *The Lancet* 392(10158): 1607–1609. [https://doi.org/10.1016/S0140-6736\(18\)32345-6](https://doi.org/10.1016/S0140-6736(18)32345-6).
- Hurd HM** (2018) *The Routledge Handbook of the Ethics of Consent*. Routledge (London).
- Information Commissioner's Office (ICO)** (2024) What Is Special Category Data?. Available at <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/lawful-basis/special-category-data/what-is-special-category-data/> (accessed 8 March 2024).
- Kelly C, Karthikesalingam A, Suleyman M, Corrado GS and King D** (2019) Key challenges for delivering clinical impact with artificial intelligence. *BMC Medicine* 17(1): 1–9. <https://doi.org/10.1186/S12916-019-1426-2>.
- Lacuna Fund** (2023) Health. Available at <https://lacunafund.org/health/> (accessed 21 November 2023).
- Lane J and Schur C** (2010) Balancing access to health data and privacy: a review of the issues and approaches for the future. *Health Services Research* 45: 1456–1467. <https://doi.org/10.1111/J.1475-6773.2010.01141.X>.
- Manson NC and O'Neill O** (2007) *Rethinking Informed Consent in Bioethics*. Cambridge University Press (Cambridge).
- McKay F, Treanor D and Hallowell N** (2023) Inalienable data: ethical imaginaries of de-identified health data ownership. *SSM- Qualitative Research in Health* 4, 100321. <https://doi.org/10.1016/j.ssmqr.2023.100321>.
- National Health Service (NHS) Confederation** (2024) Understanding Patient Data. Available at <https://understandingpatientdata.org.uk/> (accessed 8 March 2024).
- Ngusie HS et al.** (2021) Health data management practice and associated factors among health professionals working at public health facilities in resource limited settings. *Advances in Medical Education and Practice* 855–862.
- NHS England Transformation Directorate** (2023) Surveying Public Perceptions of AI – The National Strategy for AI in Health and Social Care – NHS Transformation Directorate. Available at <https://transform.england.nhs.uk/ai-lab/ai-lab-programmes/the-national-strategy-for-ai-in-health-and-social-care/surveying-public-perceptions-of-ai/> (accessed 20 November 2023).
- Nwankwo B and Sambo M** (2018) Can training of health care workers improve data management practice in health management information systems: a case study of primary health care facilities in Kaduna State, Nigeria. *Pan African Medical Journal* 30(1). <https://doi.org/10.11604/PAMJ.2018.30.289.15802>.
- Organization for Economic Cooperation and Development (OECD)** (2017) OECD Recommendation on Health Data Governance. Available at <https://www.oecd.org/els/health-systems/health-data-governance.htm> (accessed 20 November 2023).
- Owoyemi A, Osuchukwu JI, Azubuike C, ... Olaniran S** (2022) Digital solutions for community and primary health workers: lessons from implementations in Africa. *Frontiers in Digital Health* 4: 876957. <https://doi.org/10.3389/fdgh.2022.876957>.
- Pan American Health Organization (PAHO)** (2021) 8 Principles for Digital Transformation of Public Health – PAHO/WHO. Available at <https://www.paho.org/en/is4h-information-systems-health/8-principles-digital-transformation-public-health> (accessed 20 November 2023).
- Panch T, Szolovits P and Atun R** (2018) Artificial intelligence, machine learning and health systems. *Journal of Global Health* 8(2). <https://doi.org/10.7189/JOGH.08.020303>.
- Paprica PA, Melo MN de and Schull MJ** (2019) Social licence and the general public's attitudes toward research based on linked administrative health data: a qualitative study. *Canadian Medical Association Open Access Journal* 7, E40–E46. <https://doi.org/10.9778/cmajo.20180099>.
- Peiffer-Smadja N, Peiffer-Smadja N, Rawson TM, ... Holmes A** (2020) Machine learning for clinical decision support in infectious diseases: a narrative review of current applications. *Clinical Microbiology and Infection* 26(5): 584–595. <https://doi.org/10.1016/J.CMI.2019.09.009>.

- Revolutionary Government of Zanzibar** (2020) Zanzibar Digital Health Strategy 2021/21–2024/25. Available at <https://mohz.go.tz/eng/wp-content/uploads/2021/04/ZANZIBAR-DIGITAL-HEALTH-STRATEGY-2020-2025.pdf> (accessed 23 November 2023).
- Rieke N, Hancox J, ... Cardoso MJ** (2020) The future of digital health with federated learning. *npj Digital Medicine* 3(1): 1–7. <https://doi.org/10.1038/S41746-020-00323-1>.
- Sahiner B, Chen W, Samala RK and Petrick N** (2023) Data drift in medical machine learning: implications and potential remedies. *The British Journal of Radiology* 96(1150): 20220878. <https://doi.org/10.1259/bjr.20220878>.
- Sallam M** (2023) ChatGPT utility in healthcare education, research, and practice: systematic review on the promising perspectives and valid concerns. *Healthcare* 11(6): 887. <https://doi.org/10.3390/healthcare11060887>.
- Shneiderman B** (2020) Human-centered artificial intelligence: reliable, safe & trustworthy. *International Journal of Human-Computer Interaction* 36(6): 495–504. <https://doi.org/10.1080/10447318.2020.1741118>.
- Singh RP, Hom GL, Abramoff MD, Campbell JP and Chiang MF** (2020) Current challenges and barriers to real-world artificial intelligence adoption for the healthcare system, provider, and the patient. *Translational Vision Science & Technology* 9(2): 42–45. <https://doi.org/10.1167/TVST.9.2.45>.
- Sirur S, Nurse JRC and Webb H** (2018) Are we there yet? Understanding the challenges faced in complying with the General Data Protection Regulation (GDPR). In *Proceedings of the 2nd International Workshop on Multimedia Privacy and Security*. pp. 88–95. ACM (New York).
- Siyam A, Ir P, York D, ... O'Neill K** (2021) The burden of recording and reporting health data in primary health care facilities in five low- and lower-middle income countries. *BMC Health Services Research* 21(1): 1–9. <https://doi.org/10.1186/S12913-021-06652-5>.
- STANDING Together** (2023) STANdards for data Diversity, INclusivity and Generalisability. Available at <https://www.datadiversity.org/home> (accessed 21 November 2023).
- Sterckx S, Rakic V, Cockbain J and Borry P** (2016) “You hoped we would sleep walk into accepting the collection of our data”: controversies surrounding the UK care.data scheme and their wider relevance for biomedical research. *Medicine, Health Care and Philosophy* 19, 177–190. <https://doi.org/10.1007/s11019-015-9661-6>.
- Transform Health Coalition** (2023) Stronger Health Data Governance. Available at <https://transformhealthcoalition.org/health-data-governance/> (accessed 20 November 2023).
- Transform Health Coalition** (2023) Health Data Governance Principles. Available at <https://healthdatapinciples.org/> (accessed 20 November 2023).
- Utz C, Degeling M, Fahl S, Schaub F and Holz T** (2019) (Un) informed consent: studying GDPR consent notices in the field. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pp. (973–990).
- Varelius J** (2008) On the prospects of collective informed consent. *Journal of Applied Philosophy* 25(1), 35–44. <https://doi.org/10.1111/j.1468-5930.2008.00395.x>.
- Weill P and Ross JW** (2004) *IT Governance: How Top Performers Manage IT Decision Rights for Superior Results*. Harvard Business Review (Boston).
- World Bank** (2021) World Development Report 2021 – Figure 5.5. Available at <https://datacatalog.worldbank.org/search/dataset/0037980/World-Development-Report-2021—Figure-5-5> (accessed 29 April 2024).
- World Health Organization (WHO)** (2021) Global Strategy on Digital Health 2020–2025. Available at <https://iris.who.int/bitstream/handle/10665/344249/9789240020924-eng.pdf> (accessed 20 November 2023).
- World Health Organization (WHO)** (2023) WHO Health Data Governance Summit. Available at <https://www.who.int/data/events/health-data-governance-summit/introduction> (accessed 20 November 2023).