J Biomed Eng Syst Technol Low Middle Income Ctries 2026; 1(1): 76–85

Journal of Biomedical Engineering Systems and Technologies for Low- and Middle-Income Countries

Rethinking Computerized Maintenance Management Systems in LMICs

Neumann C[1 *], Douglas GP[1-3]

[1] Uptime Alliance, Nairobi, Kenya
[2] Global Health Informatics Institute, Lilongwe, Malawi
[3] University of Pittsburgh, Pittsburgh, PA, USA

*Corresponding author: Christian Neumann (christian.neumann@gmail.com)

Keywords: *Computerized Maintenance Management Systems (CMMS), Low- and Middle-Income Countries (LMICs), Artificial Intelligence (AI), Data Graveyards, Health Technology Management.*

1.0 Introduction

Although equipment graveyards remain a significant challenge in many low- and middle-income countries (LMICs),[1] the biomedical engineering community has reached broad agreement that sustainable equipment uptime depends on strong maintenance systems rather than equipment procurement alone.[2] A similar systems challenge is now becoming visible, not with the equipment itself but with the information systems used to manage it. Our subject is not the equipment graveyard but an analogous and less visible failure: the gradual deterioration of maintenance information systems through incomplete and unreliable data.

Across LMICs, investments have been made in digital systems for medical equipment inventory and maintenance management, including computerized maintenance management systems (CMMS).[3,4] In Burundi, hospitals possessed integrated CMMS functionality that was reportedly seldom used prior to reform efforts.[5] Tanzania's Medical Equipment and Infrastructure Management Information System (MEIMIS) demonstrates the growing importance of digital platforms for health technology management. However, even moderately mature national systems such as MEIMIS continue to face challenges with incomplete data entry and gaps in maintenance records.[6]

What emerges is not an equipment graveyard but what might be described as a **data graveyard**: systems that remain technically operational yet contain increasingly incomplete, outdated, or unreliable information. The distinction matters: the records are not corrupted but simply never completed, as maintenance work outpaces its documentation. The two failures share a cause: weak maintenance systems produce both poor equipment uptime and poor documentation, so the data graveyard is the counterpart of the equipment graveyard rather than its successor.

This editorial argues that recent advances in artificial intelligence (AI) may offer a practical way to address it. Rather than replacing biomedical engineers, AI can reduce the friction of capturing maintenance data, retrieving information, and supporting maintenance decisions, changing how users interact with health technology management systems.

2.0 How Data Graveyards Form

The biomedical engineering community has learned that equipment unavailability is driven not only by technical faults but also by weaknesses in maintenance systems. Preventive maintenance schedules are not followed, repair histories are not recorded, replacement parts are unavailable, and technical personnel are overwhelmed by competing demands.[3] The same pattern can occur within digital maintenance systems. Studies of CMMS implementation in LMICs have identified recurring barriers including inadequate training, inconsistent data entry, limited management support, resource constraints, and insufficient institutional ownership.[4] Rather than achieving their full potential as operational tools, information systems may gradually lose data quality and user engagement over time as maintenance records go unentered, asset inventories become outdated, and work orders remain incomplete.

Importantly, this deterioration is often gradual and largely invisible. Unlike a broken oxygen concentrator sitting in a corridor, deteriorating data quality may attract little attention until the information is needed for planning, budgeting, procurement, or policy decisions. At that point, the system may technically exist, yet no longer provide a reliable representation of reality. The challenge is therefore not merely one of software deployment. It is one of sustaining the quality, completeness, and usefulness of maintenance information over time.

3.0 The Interaction Problem

Conventional maintenance information systems were designed around computer-based data entry, structured forms and fields, and administrative maintenance workflows.[3] In contrast, the daily reality of biomedical engineers and technicians is highly mobile and task-oriented.[7] They work in wards, operating theatres, laboratories, workshops, and storage areas. They move between departments, respond to urgent breakdowns, and often manage large numbers of assets with limited resources.

Updating a maintenance record may require finding a computer, logging into a web-based system, locating the correct asset record, navigating multiple forms, and entering detailed information. When workloads are high, data entry becomes a secondary priority.[8] The maintenance work is completed, but the maintenance record is not.

Over time, this friction accumulates. The issue is not that biomedical engineers do not value documentation. Rather, the effort required to maintain digital records is often disproportionate to the immediate benefit perceived by frontline users. In practice, maintenance documentation may not be completed at the time the work is performed and may subsequently be delayed or omitted, increasing the likelihood that records are incomplete or inaccurate.[4]

The result is predictable: gradual deterioration in the completeness and accuracy of maintenance data. As records become incomplete and inaccurate, maintenance activities become harder to coordinate, monitor, and evaluate, undermining efforts to improve equipment uptime.

Framed this way, the difficulty initially appears to be one of documentation burden: recording maintenance work requires effort while offering little immediate benefit to the technician. However, on closer examination, it is better understood as a problem of interaction. Conventional maintenance management systems are static. They wait to be filled in, ask nothing, carry nothing forward from one session to the next, and offer the engineer little at the moment the work is actually being done. The resulting documentation burden reflects a fundamentally one-sided interaction: the engineer provides information to the system, but the system provides little assistance in return.

The challenges described above are not unique to maintenance information systems. Similar patterns have been observed in the design and adoption of electronic medical record (EMR) and hospital information systems in LMICs, where systems have sometimes failed to align with local workflows, organizational contexts, infrastructure, and operating environments.[9] A previously proposed "LMIC-First" framework for EMR design identified six principles that appear equally applicable to CMMS.[10] In particular, the framework argues that systems should be designed around clinical workflows and user needs rather than data collection requirements. Data should be captured as a natural byproduct of system use rather than through separate administrative processes. The framework also highlights the importance of touchscreen-first interfaces that reduce training requirements and support point-of-care use. While touchscreen interfaces offer significant advantages over traditional mouse-and-keyboard approaches, recent developments in mobile computing, speech recognition, computer vision, and generative AI create opportunities for even more natural interactions between users and information systems. Figure 1 summarizes this progression from the formation of data graveyards, through the underlying interaction problem, to the process-centric approach we propose.

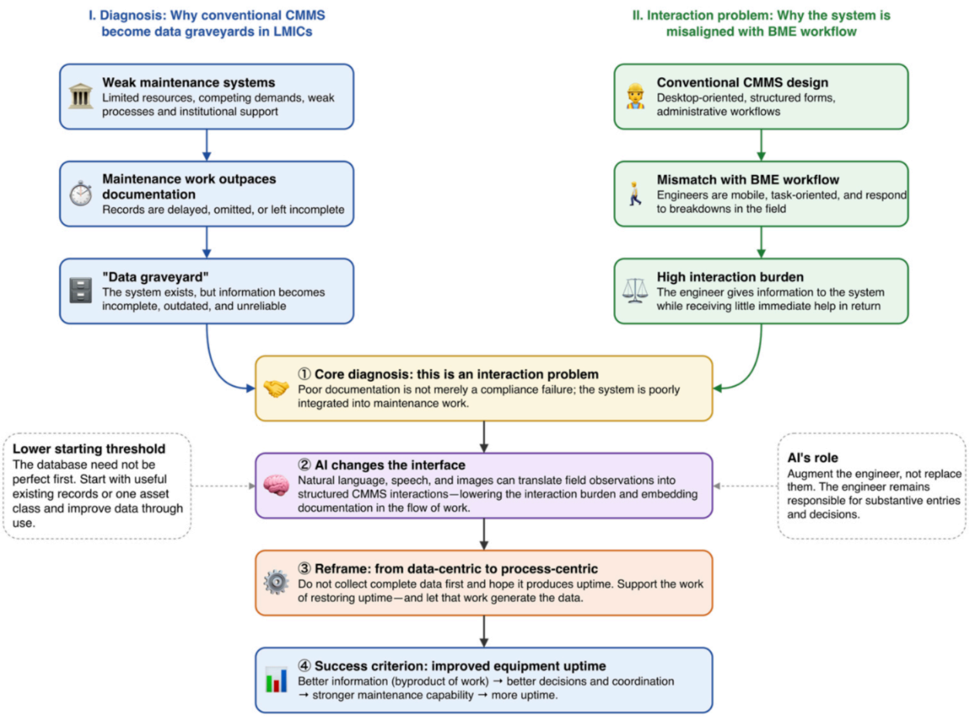

Figure 1. *From data graveyards to equipment uptime.*

4.0 Why AI May Represent a Different Paradigm

Much of the current discussion surrounding AI in healthcare focuses on diagnostic support, medical image interpretation, and clinical decision-making. While these applications are important, they may not represent the greatest opportunity for computerized maintenance management systems in LMICs.

A more immediate opportunity may lie in using AI to change how people interact with existing information systems. Natural language interfaces to databases have been developed for decades, but recent advances in large language models have substantially improved their ability to translate everyday questions into executable queries.[11] Instead of requiring users to adapt their workflow to match the software, the software can increasingly adapt to the user's workflow. These models are also multimodal, so the exchange need not be confined to typed text. A spoken note in a noisy workshop, or a photograph of a data plate or a failed component, can convey in seconds what a structured form would take much longer to capture.

In this model, AI does not replace the CMMS. It does not replace asset databases or maintenance management platforms. Instead, it functions as a conversational intermediary that translates unstructured field observations into structured maintenance records. AI becomes less of a diagnostic tool and more of a maintenance interface, helping to maintain the quality and completeness of information systems themselves.

This distinction matters for what a facility needs in place before such a system becomes useful. Deterministic software generally requires complete and well-formed data: a database query returns nothing when a record is missing, and a report is only as good as the fields behind it. A language model, by contrast, can operate on partial input and return a probable interpretation rather than an error. Where a required field is unfilled, it can request the missing value in ordinary language rather than rejecting the submission. In less structured settings, identifying what is missing is a harder problem, and current models perform inconsistently.[12,13] A facility therefore need not begin with a fully populated and cleaned database. It may instead start with the records it already holds and improve them through use, beginning with a single well-defined subset, for example oxygen concentrators or the equipment of one department, and widening as the system is used. This substantially lowers the barrier to entry. What it does require is that a person adjudicates the agent's proposals, since context that is invisible to the system, for example that a unit is already scheduled for disposal, will often determine whether a proposed action is the right one.

This perspective is consistent with Toyama's argument in *Geek Heresy* that technology is most effective when it amplifies existing human and organizational capacity rather than attempting to substitute for it.[14] In the context of health technology management, sustainable improvements in equipment uptime ultimately depend on skilled personnel, effective processes, and institutional commitment. By reducing the documentation burden, AI does not replace these foundations, it amplifies them, enabling maintenance professionals to work more effectively.

5.0 Human Augmentation in Practice

The practical implications become clearer when considered through illustrative examples.

Example A: Resolving the Interaction Bottleneck

Instead of navigating multiple web forms after completing a repair, a biomedical engineer (BME) interacts in real time with the system conversationally using a mobile phone.

BME:

"I just fixed the suction machine in the maternity ward. The motor was burned out. I replaced the motor."

Agent:

"I found one suction machine registered in the maternity ward, serial number SM-2019-0847. I'll update your open job card: fault found is burned-out motor, action taken is motor replacement, and I'll mark it as completed today. Does that look right?"

BME:

"Yes, and set the next maintenance to six months from now."

Agent:

"Done. Job card closed. The next preventive maintenance is due in six months, and I'll remind you when it's approaching."

The maintenance activity has not changed. The documentation burden has. The AI acts only as a drafting tool, while the engineer retains ultimate professional responsibility.

Example B: Asset History and Decision Support

AI may also help transform maintenance databases from passive repositories into active decision-support systems.

BME:

"Give me the history of the autoclave in the laboratory. It's failing to reach temperature again."

Agent:

"I have no autoclaves registered in the laboratory. Could you send me a photo of the data plate so I can identify the unit?"

BME:

[photograph of the data plate]

Agent:

"Thank you, that is serial number AC-772. I have attached the photo to its asset record, which had no image on file, and proposed a location correction from the pharmacy store room to the laboratory for your approval. This unit has had four job cards in the last 12 months, all related to heating elements. Total repair cost to date is US$1,200, which is approximately 60% of its original purchase value. Based on the replacement policy currently in use, this unit is a candidate for decommissioning. Do you want to proceed with the repair or initiate a replacement request?"

BME:

"Update its current location and recommend a replacement."

Agent:

"Understood. Location updated. I'll draft the replacement request and attach the maintenance history as justification for review."

In this example, AI is not performing the engineering judgment. Rather, it is reducing the effort required to access and interpret information that already exists within the system, allowing maintenance decisions to be made more quickly, consistently, and transparently in support of improved equipment uptime.

These examples point to a broader reframing. The outcome we are optimizing for shifts from complete data to increased uptime. The two are related, but they are not synonymous, and treating them as though they were is part of what produces data graveyards in the first place. Completeness is the wrong target. A database that is fully populated but never utilized does nothing for uptime. The relationship also runs in the opposite direction to the one usually assumed: rather than assembling complete data in order to improve uptime, it is the work of restoring uptime, through diagnosing the fault, ordering the part, and closing the job card, that generates the data.

This reframing substantially lowers the burden for frontline staff to engage with such a system. A conversational agent cannot compel anyone to record their work, and it should not be mistaken for a compliance mechanism; where staff are unwilling, the remedy is managerial rather than technical. We therefore make a narrower claim, which we would argue is sufficient: where the value proposition is strong enough, users adopt a system whether or not it is mandated.[15] Point-of-care electronic medical record use in Malawi, adopted at national scale during antiretroviral therapy scale-up and sustained since, is consistent with that principle.[16] The design target is not the disengaged minority but the large majority of biomedical engineers who trained for several years, want to practice their profession properly, and have not previously been given tools that make doing so easier. This is Friedman's fundamental theorem of informatics[17] applied to health technology management: the goal is not a computer that outperforms a biomedical engineer, but a biomedical engineer working with a computer who outperforms an engineer working alone.

6.0 An Opportunity for LMICs to Leapfrog

Historically, computerized health technology management systems evolved in environments characterized by ready access to personal computers, specialized software, and established institutional IT infrastructure.[18] LMICs may have an opportunity to follow a different path.

This would not be unprecedented. LMICs have demonstrated an ability to leapfrog legacy technologies. Just as many countries moved directly from limited landline infrastructure to widespread mobile phone adoption,[19] health systems may not need to replicate the desktop-centric information architectures that evolved in high-income settings. During the scale-up of HIV treatment in Malawi, a touchscreen point-of-care EMR model was adopted in place of more conventional desktop-based approaches that relied on retrospective data entry.[16,20] Rather than treating data collection as a separate activity, the system tightly coupled clinical care and information capture, allowing data to be recorded once at the point of care. The system illustrates that redesigning technology around existing workflows, rather than forcing users to adapt, can improve both usability and data quality. With smartphones already in widespread use, AI-enabled conversational interfaces may represent a similar opportunity for health technology management systems, allowing maintenance documentation to become a natural byproduct of maintenance work rather than a separate administrative task.

Importantly, while LMICs can leapfrog legacy desktop interfaces, they cannot bypass the foundational need for robust data stewardship, institutional governance, and local ownership.

7.0 Proceeding with Caution

This vision should not be mistaken for technological optimism or AI hype. AI cannot reorganize an existing data graveyard on its own, and a system with no trustworthy core to build upon will struggle. The incremental path described above lowers this barrier but does not remove it.

Significant challenges also remain regarding infrastructure, language, and governance. Despite notable improvements, rural health facilities in many African settings continue to operate with limited or unreliable internet connectivity,[21] so dependable access to cloud-hosted AI services cannot be assumed. Linguistic equity is also a hurdle. To be effective, AI models must accurately process local languages, regional dialects, and technical code-switching, rather than relying solely on high-resource languages like English.[22] Questions of data ownership and governance remain unresolved, particularly when systems are hosted outside the country in which they are used. Issues of cybersecurity, privacy, interoperability, and regulatory oversight will also require careful consideration.[23]

Technical challenges should not be underestimated. Unlike traditional software systems, AI models are inherently probabilistic rather than deterministic. As a result, they may generate inaccurate or misleading outputs, creating new challenges for system validation, testing, and quality assurance. Human oversight, appropriate safeguards, and clearly defined operational boundaries will therefore remain essential.[23] To mitigate the safety and liability risks of AI hallucinations, the model must remain an interface to the record rather than an author of it. Every substantive entry originates from the engineer and is confirmed before it is committed, with the engineer retaining ultimate professional responsibility.

Long-term sustainability is another consideration. The future cost of operating large language models remains uncertain, and health systems must avoid dependence on proprietary platforms they cannot maintain locally. Efficiency trends work in favor of adoption: performance is no longer dependent on sheer scale. Recently, the smallest model able to pass a widely used test of general knowledge was around 140 times smaller than the model that first cleared the same mark just over two years earlier.[24] Models at this scale run outside specialized computing facilities,[25] making self-managed deployment more viable for institutions without large infrastructure.

Data quality does not sustain itself. Asset inventories do not update themselves. Maintenance histories do not remain accurate without continual effort. Like medical devices, information systems require maintenance.

8.0 Conclusion

The equipment graveyard taught the biomedical engineering community that uptime depends on maintenance systems, not procurement alone. The data graveyard suggests a parallel lesson. Information quality depends on how well data capture is built into the work itself, not on how reliably it is reconstructed afterwards.

If we reframe our view of these systems from data-centric to process-centric, documentation can become a byproduct of maintenance rather than a task competing with it. A partially complete record that is acted upon contributes more to uptime than an exhaustive one that is never consulted. It also lowers the barrier to starting: a facility need not clean its database first, but can begin with a single asset class and improve the records through use.

The next frontier of AI in biomedical engineering may therefore be surprisingly mundane. Rather than diagnosing disease or replacing technicians, its value will lie in capturing, organizing, and retrieving maintenance information, lowering the interaction burden that limits routine use of these systems. AI removes administrative friction, but it cannot substitute for institutional leadership, maintenance budgets, and a culture of accountability. The question is whether these tools can make information systems usable enough to become part of everyday practice. Ultimately, the measure of success is equipment uptime, not data completeness.

Authors' Contributions

CN and GPD contributed to the conceptualization and investigation of the work. CN prepared the first draft. GPD provided critical input on the second draft. Both authors contributed to reviewing and editing subsequent versions and approved the final version of the manuscript.

Declaration of Interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Declaration of Generative AI Assistance

The authors used large language models to assist with language editing and revision of the manuscript. A large language model was also used to assist in organizing the authors' argument into the conceptual representation shown in Figure 1. The authors independently verified all cited sources, reviewed and edited all AI-assisted material, and take full responsibility for the content of the manuscript.

References

1. Asma E, Heenan M, Banda G, Kirby RP, Mangwiro L, Acemyan CZ, et al. Avoid equipment graveyards: rigorous process to improve identification and procurement of effective, affordable, and usable newborn devices in low-resource hospital settings. BMC Pediatr. 2023 Nov 15;23(S2):569. doi:10.1186/s12887-023-04362-x
2. World Health Organization. Medical Equipment Maintenance Programme Overview [Internet]. Geneva: World Health Organization; 2011 [cited 2026 Aug 15]. (WHO Medical device technical series). Available from: https://iris.who.int/handle/10665/44587
3. World Health Organization. Inventory and maintenance management information system for medical devices [Internet]. Geneva: World Health Organization; 2025 [cited 2026 Aug 15]. (WHO medical device technical series). Available from: https://www.who.int/publications/i/item/9789240111257
4. Cohen D, Visnjic N, Akaateba D, Hadfield K. Narrative literature review of facilitators and barriers to implementing computerized maintenance management systems in low-middle-income countries. Health Technol. 2023 Jun;13(3):373–8. doi:10.1007/s12553-023-00743-5
5. Beniacoub F, Ntwari F, Niyonkuru JP, Nyssen M, Van Bastelaere S. Evaluating a computerized maintenance management system in a low resource setting. Health Technol. 2021 May;11(3):655–61. doi:10.1007/s12553-021-00524-y
6. Abdallah AK, Haule S, Werlein R, Mvanga V, Delcroix P, Saric J, et al. Medical device management reform,

United Republic of Tanzania. Bull World Health Organ. 2024;102(9):665–73. doi:10.2471/BLT.23.290636
7. Hillebrecht M, Schmidt C, Saptoka BP, Riha J, Nachtnebel M, Bärnighausen T. Maintenance versus replacement of medical equipment: a cost-minimization analysis among district hospitals in Nepal. BMC Health Serv Res. 2022 Aug 12;22(1):1023. doi:10.1186/s12913-022-08392-6
8. Inagaki D, Nakahara S, Chung U, Shimaoka M, Shoji K. Need for improvements in medical device management in low- and middle-income countries: applying learnings from Japan's experience. JMA J. 2023;6(2):188–91. doi:10.31662/jmaj.2022-0089
9. Littlejohns P, Wyatt JC, Garvican L. Evaluating computerised health information systems: hard lessons still to be learnt. BMJ. 2003 Apr 19;326(7394):860–3. doi:10.1136/bmj.326.7394.860
10. Neumann C, Dunbar EL, Espino JU, Mtonga TM, Douglas GP. A LMIC-first manifesto to developing electronic medical record systems. J Health Inform Afr. 2021 Jun 12;7(2):1–3. doi:10.12856/JHIA-2020-v7-i2-300
11. Liu M, Wang X, Xu J, Yi W, Wolfson O. A systematic review of natural language interfaces for databases. Front Comput Sci. 2026;20(11):2011623. doi:10.1007/s11704-025-50592-w
12. Li BZ, Kim B, Wang Z. QuestBench: Can LLMs ask the right question to acquire information in reasoning tasks? [Internet]. arXiv; 2025 [cited 2026 Aug 12]. Available from: https://arxiv.org/abs/2503.22674. doi:10.48550/arXiv.2503.22674
13. Zhou Z, Feng X, Zhu Z, Yao J, Koyejo S, Han B. From passive to active reasoning: can large language models ask the right questions under incomplete information? [Internet]. arXiv; 2025 [cited 2026 Aug 14]. Available from: https://arxiv.org/abs/2506.08295. doi:10.48550/arXiv.2506.08295
14. Toyama K. Geek heresy: rescuing social change from the cult of technology. New York: PublicAffairs; 2015.
15. Richards J, Douglas G, Fraser HSF. Perspectives on global public health informatics. In: Magnuson JA, Fu P, editors. Public health informatics and information systems [Internet]. London: Springer; 2014 [cited 2026 Aug 14]. p. 619–44. Available from: https://link.springer.com/10.1007/978-1-4471-4237-9_31. doi:10.1007/978-1-4471-4237-9_31
16. Douglas GP, Gadabu OJ, Joukes S, Mumba S, McKay MV, Ben-Smith A, et al. Using touchscreen electronic medical record systems to support and monitor national scale-up of antiretroviral therapy in Malawi. PLoS Med. 2010 Aug 10;7(8):e1000319. doi:10.1371/journal.pmed.1000319
17. Friedman CP. A "fundamental theorem" of biomedical informatics. J Am Med Inform Assoc. 2009 Mar 1;16(2):169–70. doi:10.1197/jamia.M3092
18. Cohen T, Baretich MF, Gentles WM. Computerized maintenance management systems. In: Clinical Engineering Handbook [Internet]. Amsterdam: Elsevier; 2020 [cited 2026 Aug 11]. p. 208–18. Available from: https://linkinghub.elsevier.com/retrieve/pii/B9780128134672000341. doi:10.1016/B978-0-12-813467-2.00034-1
19. James J. Leapfrogging in mobile telephony: a measure for comparing country performance. Technol Forecast Soc Change. 2009;76(7):991–8. doi:10.1016/j.techfore.2008.09.002
20. Derksen L, McGahan AM, Pongeluppe LS. The lifesaving impact of electronic

medical records for HIV patients. Rev Econ Stat. 2025 Oct 29;1–46. doi:10.1162/REST.a.1617
21. Bataliack S, Ebongue M, Karamagi H, Janauschek L. Health data digitalization in Africa: unlocking the potential [Policy Brief] [Internet]. Brazzaville: World Health Organization Regional Office for Africa; 2024 [cited 2026 Aug 12]. (AHOP Policy Briefs). Available from: https://iris.who.int/items/92a48c20-aec1-49cd-8737-39acc33ad2fb
22. Olaleye K, Oncevay A, Sibue M, Zondi N, Terblanche M, Mapikitla S, et al. AfroCS-xs: creating a compact, high-quality, human-validated code-switched dataset for African languages. In: Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers) [Internet]. Vienna, Austria: Association for Computational Linguistics; 2025 [cited 2026 Aug 12]. p. 33391–410. Available from: https://aclanthology.org/2025.acl-long.1601. doi:10.18653/v1/2025.acl-long.1601
23. World Health Organization. Ethics and governance of artificial intelligence for health: WHO guidance [Internet]. Geneva: World Health Organization; 2021 [cited 2026 Aug 15]. Available from: https://www.who.int/publications/i/item/9789240029200
24. Maslej N, Fattorini L, Perrault R, Gil Y, Parli V, Kariuki N, et al. Artificial intelligence index report 2025 [Internet]. arXiv; 2025 [cited 2026 Aug 15]. Available from: https://arxiv.org/abs/2504.07139. doi:10.48550/arXiv.2504.07139
25. Abdin M, Aneja J, Awadalla H, Awadallah A, Awan AA, Bach N, et al. Phi-3 technical report: a highly capable language model locally on your phone [Internet]. arXiv; 2024 [cited 2026 Aug 14]. Available from: https://arxiv.org/abs/2404.14219. doi:10.48550/arXiv.2404.14219