BMJ Health & Care Informatics

Scaling equitable artificial intelligence in healthcare with machine learning operations

Madelena Y Ng ,[1,2] Alexey Youssef ,[3,4] Malvika Pillai ,[5,1] Vaibhavi Shah ,[1] Tina Hernandez-Boussard [1,2]

To cite: Ng MY, Youssef A, Pillai M, *et al*. Scaling equitable artificial intelligence in healthcare with machine learning operations. *BMJ Health Care Inform* 2024;**31**:e101101. doi:10.1136/bmjhci-2024-101101

Received 12 April 2024
Accepted 25 September 2024

[1]Department of Medicine (Biomedical Informatics), Stanford University, Stanford, California, USA
[2]Department of Biomedical Data Science, Stanford University, Stanford, California, USA
[3]Department of Engineering Science, Oxford University, Oxford, UK
[4]Department of Bioengineering, Stanford University, Stanford, California, USA
[5]VA Palo Alto Health Care System, Palo Alto, California, USA

Correspondence to
Dr Madelena Y Ng;
madelena@stanford.edu

Machine learning operations (MLOps), a discipline concerned with the production, monitoring and maintenance of artificial intelligence (AI) and machine learning (ML) models at scale, applied in healthcare can facilitate the transition of AI/ML-enabled healthcare tools from research to sustainable deployment.[1–3] Adherence to MLOps best practices can address persistent challenges with AI/ML tools deployed into clinical workflows where models often struggle with generalisability, integration and robustness. As AI regulations continue to evolve such as the Department of Health and Human Services Office of Civil Rights final ruling that requires healthcare providers to ensure their AI/ML tools do not discriminate,[4] it becomes increasingly essential for MLOps in healthcare to prioritise health equity.

MLOps in healthcare contextualises core principles to serve the needs of AI/ML deployment at healthcare organisations.[1–3] We highlight key principles and considerations to scale equitable health AI/ML model deployments and establish accountability measures[5] that systematically quell health inequities and comply with AI/ML regulations (table 1).

One MLOps in healthcare principle is to optimise the clinical workflow integration of health AI/ML tools for care delivery. Additionally, ensuring models are accessible to all patient populations and adaptable to different clinical settings is crucial. A suitable functionality that prioritises health equity is to conduct clinical workflow analyses to identify key stakeholders, processes and resources required to execute AI/ML equitably for patient care. The paths identified inform the development of workflow orchestration components that serve as the foundation of MLOps pipelines at the healthcare organisation.

Another principle is that diverse stakeholders can collaborate to address the operational needs of health AI/ML tools. Effective operations require the coordination of knowledge and efforts across an organisation. Functionalities that prioritise health equity in collaboration include the documentation of health equity assessments[6] and the creation of an equitable AI expert panel (eg, healthcare providers, researchers, patients, ethics and compliance officers, diversity, equity, and inclusion officers, privacy officers) to inform decision-making. These components will enhance transparency and help justify the performance, safety and equity thresholds set for deployed AI/ML tools in MLOps healthcare infrastructure.

A defining principle is the continuous maintenance, monitoring and support of health AI/ML tools over time to ensure the most optimal version is in production for clinical impact. Functionalities that promote health equity include reinforcing model monitoring components with the detection of fairness drifts in AI/ML tools. This will require setting triggers for model recalibration, updating or retraining based on key fairness metrics (eg, equalised odds, equal opportunity, demographic parity)[7] prior to deployment. Continuously monitoring model performance, safety and equity over time and the ability to reconcile conflicting deviations ensures holistically calibrated AI/ML models reach patients. Furthermore, this component of MLOps in healthcare infrastructure prepares organisations to meet regulatory criteria (eg, predetermined change control plans) and can expedite the Food and Drug Administration-approval of home-grown AI/ML models for broader clinical use.[8]

MLOps in healthcare also follows the principle of automation regarding the continuous delivery (CD), continuous integration

Table 1 Considerations to prioritise health equity in MLOps healthcare infrastructure

MLOps in healthcare principles	Functional requirements that prioritise health equity	Healthcare AI/ML deployment challenges addressed
Clinical workflow integration of health AI/ML models	Conduct clinical workflow analyses to identify paths on how to deploy AI/ML equitably.	Coordination and efficient usage of resources to identify and mitigate AI/ML biases and improve the embedding of AL/ML models into diverse clinical settings.
Collaboration and communication across diverse stakeholders	Documentation of health equity assessments for AI/ML tools and the creation of an equitable AI expert panel.	Transparency about how AI/ML tools are evaluated and determined to best serve patients considering equity, safety and performance.
Continuous maintenance, monitoring and support of health AI/ML models	Incorporate detection of fairness drifts in AI/ML tools during model monitoring and set triggers based on key fairness metrics in advance of deployment.	Ensures holistically calibrated AI/ML models reach patients and organisations are prepared to meet regulatory opportunities.
Automation of continuous delivery, continuous integration and continuous deployment of health AI/ML models	Automate bias and fairness checks within the deployment pipeline and ensure feature stores include features critical to advancing health equity such as social determinants of health.	Alleviate associated AI/ML operational burdens (eg, costs, time, quality) for an organisation and ensure consistency in bias and fairness checks.
Compliance with evolving regulations on health AI/ML and patient safety	Expand automated quality checks and assurances to screen AI/ML tools for bias and fairness before and after deployment.	Prevents the deployment or use of AI/ML models that pose excessive risks to equity, patients and the organisation.

AI, artificial intelligence; ML, machine learning; MLOps, machine learning operations.

(CI) and continuous deployment of AI/ML tools. This principle requires the setup of CD/CI components that automate steps to enable rapid iteration and model optimisation on receiving feedback. The automation of manual processes can alleviate associated AI/ML operational burdens (eg, costs, time, quality) for an organisation and ensure consistency in bias and fairness checks. To promote health equity through this component, one functional requirement is that feature stores include features considered critical to advancing health equity such as social determinants of health. The ability to standardise health equity feature processing pipelines that feed into feature stores will enable engineers to scale and ensure models are 'fit-for-equity'.

Finally, an underlying principle is that MLOps in healthcare require compliance with evolving regulations on health AI/ML and patient safety. Organisations must keep pace with changing standards and modify technical guardrails accordingly to provide proper oversight while ensuring AI/ML models do not inadvertently harm vulnerable populations. One functional requirement is to expand automated quality and assurance checks to screen AI/ML tools for bias and fairness before deployment and monitor for discrimination and jailbreaking after deployment. Those unable to meet minimum thresholds of quality in terms of performance, safety and equity are considered not compliant with regulation and pose an excessive risk to worsening health inequities or organisational reputation.

With the surge of health AI innovation, forethought on how to integrate functional health equity accountability measures into MLOps healthcare infrastructure can catalyse the adoption and scale of equitable AI/ML-enabled tools for broad clinical impact. The path forward in healthcare requires recognition, investment and adoption of robust MLOps pipelines to advance operational excellence for equitable AI.

X Malvika Pillai @MalvikaPillai

Contributors MYN and AY contributed equally to this paper. Study concept and design: MYN, AY and TH-B. Analysis and interpretation of data: All authors (MYN, AY, MP, VS and TH-B). Writing—original draft preparation: MYN and AY. Writing—critical revision of the manuscript for important intellectual content: All authors (MYN, AY, MP, VS and TH-B). All authors read and approved the final manuscript. TH-B is the guarantor of this work.

Funding Research reported in this publication was supported by the National Center for Advancing Translational Sciences under Award Number UL1TR003142, Agency for Healthcare Research and Quality (AHRQ) under Award Number R01HS027434 and the National Library of Medicine under Award Number T15LM007033. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health or AHRQ.

Competing interests None declared.

Patient consent for publication Not applicable.

Ethics approval Not applicable.

Provenance and peer review Not commissioned; externally peer reviewed.

ORCID iDs
Madelena Y Ng http://orcid.org/0000-0003-3824-9349
Alexey Youssef http://orcid.org/0000-0003-0621-1522

Malvika Pillai http://orcid.org/0000-0001-8739-189X
Vaibhavi Shah http://orcid.org/0000-0001-5576-9878
Tina Hernandez-Boussard http://orcid.org/0000-0001-6553-3455

REFERENCES

1 Khattak FK, Subasri V, Krishnan A, *et al*. MLHOps: machine learning for healthcare operations (pre print) [arXiv]. 2023.
2 Rajagopal A, Ayanian S, Ryu AJ, *et al*. Machine Learning Operations in Health Care: A Scoping Review. *Mayo Clin Proc Dig Health* 2024;2:421–37.
3 Kreuzberger D, Kühl N, Hirschl S. Machine Learning Operations (MLOps): Overview, Definition, and Architecture. *IEEE Access* 2023;11:31866–79.
4 U.S. Department of Health and Human Services (HHS) Office for Civil Rights (OCR). HHS issues new rule to strengthen nondiscrimination protections and advance civil rights in health care. 2024. Available: https://www.hhs.gov/about/news/2024/04/26/hhs-issues-new-rule-strengthen-nondiscrimination-protections-advance-civil-rights-health-care.html
5 Chin MH, Afsar-Manesh N, Bierman AS, *et al*. Guiding Principles to Address the Impact of Algorithm Bias on Racial and Ethnic Disparities in Health and Health Care. *JAMA Netw Open* 2023;6:e2345050.
6 Kim JY, Hasan A, Kellogg KC, *et al*. Development and preliminary testing of Health Equity Across the AI Lifecycle (HEAAL): A framework for healthcare delivery organizations to mitigate the risk of AI solutions worsening health inequities. *PLOS Dig Health* 2024;3:e0000390.
7 Castelnovo A, Crupi R, Greco G, *et al*. A clarification of the nuances in the fairness metrics landscape. *Sci Rep* 2022;12:4209.
8 U.S. Food and Drug Administration (FDA). Predetermined change control plans for machine learning-enabled medical devices: guiding principles. 2023 Available: https://www.fda.gov/medical-devices/software-medical-device-samd/predetermined-change-control-plans-machine-learning-enabled-medical-devices-guiding-principles