


Keeping a human in the loop: Managing the ethics of AI in medicine

October 20 2023



Credit: Pixabay/CC0 Public Domain

Artificial intelligence (AI)—of ChatGPT fame—is increasingly used in medicine to improve diagnosis and treatment of diseases, and to avoid

unnecessary screening for patients. But AI medical devices could also harm patients and worsen health inequities if they are not designed, tested, and used with care, according to an international task force that included a University of Rochester Medical Center bioethicist.

Jonathan Herington, Ph.D., was a member of the AI Task Force of the Society for Nuclear Medicine and Medical Imaging, which laid out recommendations on how to [ethically develop](#) and [use AI medical devices](#) in two papers published in the *Journal of Nuclear Medicine*.

In short, the [task force](#) called for increased transparency about the accuracy and limits of AI and outlined ways to ensure all people have access to AI [medical devices](#) that work for them—regardless of their race, ethnicity, gender, or wealth.

While the burden of proper design and testing falls to AI developers, [health care providers](#) are ultimately responsible for properly using AI and shouldn't rely too heavily on AI predictions when making patient care decisions.

"There should always be a human in the loop," said Herington, who is assistant professor of Health Humanities and Bioethics at URMC and was one of three bioethicists added to the task force in 2021. "Clinicians should use AI as an input into their own decision making, rather than replacing their decision making."

This requires that doctors truly understand how a given AI [medical device](#) is intended to be used, how well it performs at that task, and any limitations—and they must pass that knowledge on to their patients. Doctors must weigh the relative risks of false positives versus false negatives for a given situation, all while taking structural inequities into account.

When using an AI system to identify probable tumors in PET scans, for example, health care providers must know how well the system performs at identifying this specific type of tumor in patients of the same sex, race, ethnicity, etc., as the patient in question.

"What that means for the developers of these systems is that they need to be very transparent," said Herington.

According to the task force, it's up to the AI developers to make [accurate information](#) about their medical device's intended use, clinical performance, and limitations readily available to users.

One way they recommend doing that is to build alerts right into the device or system that informs users about the degree of uncertainty of the AI's predictions. That might look like heat maps on cancer scans that show whether areas are more or less likely to be cancerous.

To minimize that uncertainty, developers must carefully define the data they use to train and test their AI models, and should use clinically relevant criteria to evaluate the model's performance. It's not enough to simply validate algorithms used by a device or system.

AI medical devices should be tested in so-called "silent trials," meaning their performance would be evaluated by researchers on real patients in real time, but their predictions would not be available to the [health care provider](#) or applied to clinical decision making.

Developers should also design AI models to be useful and accurate in all contexts in which they will be deployed.

"A concern is that these high-tech, expensive systems would be deployed in really high-resource hospitals, and improve outcomes for relatively well-advantaged patients, while patients in under-resourced or rural

hospitals wouldn't have access to them—or would have access to systems that make their care worse because they weren't designed for them," said Herington.

Currently, AI medical devices are being trained on datasets in which Latino and Black patients are underrepresented, meaning the devices are less likely to make accurate predictions for patients from these groups.

In order to avoid deepening [health inequities](#), developers must ensure their AI models are calibrated for all racial and gender groups by training them with datasets that represent all of the populations the medical device or system will ultimately serve.

Though these recommendations were developed with a focus on [nuclear medicine](#) and [medical imaging](#), Herington believes they can and should be applied to AI medical devices broadly.

"The systems are becoming ever more powerful all the time and the landscape is shifting really quickly," said Herington. "We have a rapidly closing window to solidify our ethical and regulatory framework around these things."

More information: Jonathan Herington et al, Ethical Considerations for Artificial Intelligence in Medical Imaging: Data Collection, Development, and Evaluation, *Journal of Nuclear Medicine* (2023). [DOI: 10.2967/jnumed.123.266080](https://doi.org/10.2967/jnumed.123.266080) Jonathan

Herington et al, Ethical Considerations for Artificial Intelligence in Medical Imaging: Deployment and Governance, *Journal of Nuclear Medicine* (2023). [DOI: 10.2967/jnumed.123.266110](https://doi.org/10.2967/jnumed.123.266110)

Provided by University of Rochester Medical Center

Citation: Keeping a human in the loop: Managing the ethics of AI in medicine (2023, October 20) retrieved 28 April 2024 from <https://medicalxpress.com/news/2023-10-human-loop-ethics-ai-medicine.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.