



← Home ← Insights ← Industries ← Healthcare

ARTICLE

May 06, 2024

The Ethics of Generative AI in Healthcare: Addressing Algorithmic Bias

Artificial intelligence (AI) represents a new frontier for healthcare, offering many opportunities to improve health outcomes across patient populations.

At the same time, however, it comes with serious risks stemming from its potential for bias. In order to adopt AI safely and successfully, healthcare organizations must maintain a patient-first mentality, understanding the risks that AI bias presents and taking proper steps to mitigate them.

The AI Landscape in Healthcare

Traditionally, healthcare has been slow to adopt new technology, but healthcare leaders report they're moving quickly on AI. In fact, [BDO's 2024 Healthcare CFO Outlook Survey](#) found that 98% of organizations surveyed are already piloting generative AI programs. At present, these programs are most often used to lessen the administrative load on healthcare professionals and help

alleviate clinician burnout. For example, many healthcare professionals are making use of note-taking and ambient-listening programs to securely and accurately log notes during consultations, subject to evolving privacy and security regulations. However, AI is also being piloted for use in patient care, including generating treatment plans that are subject to human review, based on symptoms and available patient history.

As AI tools become more sophisticated and their large, generalized data pools are improved through constant maintenance and good data governance, these tools will be used to create more and more personalized treatment plans. Specifically, generative AI tools will be able to compare a patient's unique circumstances — their symptoms, history, medications, and applicable Social Determinants of Health (SDoH), such as race, class, or gender — against increasingly specific and similar data groups to better tailor treatment plans.

Generative AI can be used to develop a highly detailed patient history that can help physicians rule out illnesses and identify potential diagnoses more quickly. It can also map out data-driven treatment options, which can help improve preventive care and, in the long term, mitigate expenses related to chronic care.

Ultimately, AI tools have the power to transform healthcare. However, as with any technology, AI brings risks that healthcare leaders must proactively address to protect their patients and their organization.

Confronting AI Bias

When adopting AI, healthcare organizations must take active steps to mitigate the considerable risk presented by [AI bias](#). Access to care and patient safety are already areas of concern for marginalized populations, and real-world biases can be carried over into the digital treatment space if an AI model is trained on biased data. For this reason, healthcare organizations should consistently monitor their data to avoid amplifying healthcare disparities.

Human oversight and avoiding overreliance by clinicians is also essential when it comes to the ethical use of AI, because no single algorithm can account for all SDoH and address every patient population. For example, drawing on a data pool of primarily female patients might generate less effective or even harmful results when applied to male patients. For that reason, there must be a human in the loop to identify potential instances of bias whenever AI is used.

However, the question of how patient data is collected, maintained, and used further complicates how healthcare providers should use AI. On the one hand, strong AI tools generally rely on the largest possible data pools to drive their insights, so it can be tempting for healthcare organizations to engage with publicly available and data-rich large language models (LLMs) built by third parties.

However, there is always the possibility that third parties feed their LLM with poorly curated data or questionable research studies that could lead their AI model to create biased clinical decision recommendations. Clinicians should never take the diagnostics recommendations of AI at face-value and should always flag when AI-generated clinical decision support seems to stem from bias.

When healthcare organizations outsource the governance and maintenance of their AI model to companies unfamiliar with the healthcare industry, it becomes more difficult to implement the necessary human oversight to identify bias generated by AI tools. Maintaining strong data governance, and training clinicians to identify AI bias is essential to help providers avoid making clinical decisions based on a flawed understanding of AI generated outputs.

On the other hand, a healthcare organization's ability to train AI tools using its own data is limited by Health Insurance Portability and Accountability Act (HIPAA) [security rules](#) and requires significant investment and infrastructure. Many healthcare organizations may find the cost of building and maintaining that infrastructure challenging or unachievable in the current economic climate.

Compliance & Auditing

In early February, the Centers for Medicare and Medicaid Services (CMS) [issued a memo](#) stipulating that health insurance companies cannot use AI or algorithms to determine coverage or deny members on Medicare Advantage (MA) plans without human oversight. Additional regulatory scrutiny has already been enacted in the European Union through the [EU AI Act](#), which would enshrine that AI in healthcare is always deployed in the best interests of patients. Similar regulatory action is likely on the horizon in the U.S. to regulate the use of treatment plans generated with the support of AI.

AI's future in healthcare depends on organizations [keeping a human in the loop](#) to maintain

responsible usage. Over time, these tools will only become more prevalent in clinical and teaching settings. Putting patients first by instituting a responsible, human-driven AI program now can help organizations get ahead of inevitable regulatory changes, which are likely to focus on requiring human oversight of AI use. Without upfront action, these regulatory changes could introduce new compliance concerns that might prove disruptive to healthcare and expensive for providers.

The National Institute of Health has outlined a strategy for addressing bias issues when deploying AI tools in a healthcare setting called the “[Total Product Lifecycle \(TPLC\)](#)” model. Rather than approaching potential bias as a problem to identify and to fix, the TPLC model helps healthcare providers prevent bias during every stage of AI implementation, from Conceptualization to Design, Development, Validation, Access, and Monitoring.

Monitoring and evaluation of AI programs, once implemented, are equally as important. A recent panel held at Yale University outlined a series of principles for [mitigating algorithmic bias](#) in healthcare and highlighted the need for ongoing evaluation of algorithm performance and outcomes. Healthcare organizations should give careful thought to who is best positioned to lead these evaluations. Additionally, healthcare organizations need more than just internal controls to maintain an ethical AI program — organizations should plan to include AI as part of their annual external audit programs, to improve controls across their enterprise.

Five Steps to Responsible AI

There is no “one size fits all” approach to creating a responsible AI program in healthcare. But no matter the size or type of a business, there are five steps all organizations should take to set the foundation for an ethical AI program:

1. Educate

The first step is to understand what AI is, how it can benefit the organization, and the risks that it presents.

2. Find Relevant Use Cases

The real power of AI emerges when tailored to an organization's specific needs. A good first use case for AI is one that represents a manageable risk to the business with the most immediate benefit. In many healthcare organizations, a common first use case is note taking and charting, that can reduce the administrative burden on clinicians.

3. Prepare and Build

Just like a physical building, the strength of an AI initiative lies in its foundation, and the foundation of AI is in high-quality, diverse data. Once that data is in place and supported by a robust data management environment that prioritizes data security and data governance, healthcare organizations need to implement controls to maintain compliance with data privacy standards and to safeguard sensitive and personal data. Only after these two steps are completed can organizations build a responsible AI program.

4. Enable and Adopt

A tool is only as good as its user. Integrating AI into workflows requires a cultural shift. Providing training and resources is key, and most importantly, everyone should understand the 'why' behind the AI. Healthcare organizations may encounter friction internally, especially from already overburdened clinicians with little time to spend on training. However, emphasizing how clinicians can benefit from the technology, supported by use cases, especially in terms of lightening their workload and supporting their high quality of care, may help encourage clinicians to actively participate in the training.

5. Go & Grow

Finally, once the groundwork is laid, the AI can be launched. But the AI journey doesn't end here; as an organization evolves, so will its AI needs. Regularly revisiting, refining, and iterating will help AI systems stay relevant and continue delivering value. It's also crucial to constantly revisit the program, incorporating AI into yearly external auditing to verify that it is working as intended and providing high-quality outputs that serve patient needs.

Want to Know More About How AI Is Already Changing Healthcare?

Read BDO's insight on AI and predictive staffing models.

ARTICLE

**Powering Predictive Staffing in
Healthcare with AI**

