

Stratospheric safety standards: How aviation could steer regulation of AI in health

January 17 2024, by Alex Ouyang



Credit: Pixabay from Pexels

What is the likelihood of dying in a plane crash? According to a 2022

report released by the International Air Transport Association, the industry fatality risk is 0.11. In other words, on average, a person would need to take a flight every day for 25,214 years to have a 100 percent chance of experiencing a fatal accident. Long touted as one of the safest modes of transportation, the highly regulated aviation industry has MIT scientists thinking that it may hold the key to regulating artificial intelligence in health care.

Marzyeh Ghassemi, an assistant professor at the MIT Department of Electrical Engineering and Computer Science (EECS) and Institute of Medical Engineering Sciences, and Julie Shah, an H.N. Slater Professor of Aeronautics and Astronautics at MIT, share an interest in the challenges of transparency in AI models. After chatting in early 2023, they realized that aviation could serve as a model to ensure that marginalized patients are not harmed by biased AI models.

Ghassemi—who is also a principal investigator at the MIT Abdul Latif Jameel Clinic for Machine Learning in Health (Jameel Clinic) and the Computer Science and Artificial Intelligence Laboratory (CSAIL)—and Shah then recruited a cross-disciplinary team of researchers, attorneys, and policy analysts across MIT, Stanford University, the Federation of American Scientists, Emory University, University of Adelaide, Microsoft, and the University of California San Francisco to kick off a research project, the results of which were recently accepted to the [Equity and Access in Algorithms, Mechanisms and Optimization Conference](#).

"I think I can speak for both Marzyeh and myself when I say that we're really excited to see kind of excitement around AI starting to come about in society," says first author Elizabeth Bondi-Kelly, now an assistant professor of EECS at the University of Michigan who was a postdoc in Ghassemi's lab when the project began. "But we're also a little bit cautious and want to try to make sure that it's possible we can have

frameworks in place to manage potential risks as these deployments start to happen, so we were looking for inspiration for ways to try to facilitate that."

AI in [health](#) today bears a resemblance to where the aviation industry was a century ago, says co-author Lindsay Sanneman, a Ph.D. student in the Department of Aeronautics and Astronautics at MIT. Though the 1920s were known as "the Golden Age of Aviation," [fatal accidents](#) were "disturbingly numerous," according to the Mackinac Center for Public Policy.

Jeff Marcus, the current chief of the National Transportation Safety Board (NTSB) Safety Recommendations Division, recently published a National Aviation Month blog post noting that while a number of fatal accidents occurred in the 1920s, 1929 remains the "worst year on record" for the most fatal aviation accidents in history, with 51 reported accidents. By today's standards that would be 7,000 accidents per year, or 20 per day. In response to the high number of fatal accidents in the 1920s, President Calvin Coolidge passed landmark legislation in 1926 known as the Air Commerce Act, which would regulate air travel via the Department of Commerce.

But the parallels do not stop there—aviation's subsequent path into automation is similar to AI's. AI explainability has been a contentious topic given AI's notorious "black box" problem, which has AI researchers debating how much an AI model must "explain" its result to the user before potentially biasing them to blindly follow the model's guidance.

"In the 1970s there was an increasing amount of automation ... autopilot systems that take care of warning pilots about risks," Sanneman adds. "There were some growing pains as automation entered the aviation space in terms of human interaction with the autonomous

system—potential confusion that arises when the pilot doesn't have keen awareness about what the automation is doing."

Today, becoming a commercial airline captain requires 1,500 hours of logged flight time along with instrument trainings. According to the researchers' paper, this rigorous and comprehensive process takes approximately 15 years, including a bachelor's degree and co-piloting. Researchers believe the success of extensive pilot training could be a potential model for training [medical doctors](#) on using AI tools in clinical settings.

The paper also proposes encouraging reports of unsafe health AI tools in the way the Federal Aviation Agency (FAA) does for pilots—via "limited immunity," which allows pilots to retain their license after doing something unsafe, as long as it was unintentional.

According to a [2023 report](#) published by the World Health Organization, on average, one in every 10 patients is harmed by an adverse event (i.e., "medical errors") while receiving hospital care in high-income countries.

Yet in current health care practice, clinicians and [health care workers](#) often fear reporting medical errors, not only because of concerns related to guilt and self-criticism, but also due to negative consequences that emphasize the punishment of individuals, such as a revoked medical license, rather than reforming the system that made medical error more likely to occur.

"In health, when the hammer misses, patients suffer," wrote Ghassemi in a recent [comment published](#) in *Nature Human Behavior*. "This reality presents an unacceptable ethical risk for medical AI communities who are already grappling with complex care issues, staffing shortages, and overburdened systems."

Grace Wickerson, co-author and health equity policy manager at the Federation of American Scientists, sees this new paper as a critical addition to a broader governance framework that is not yet in place. "I think there's a lot that we can do with existing government authority," they say.

"There's different ways that Medicare and Medicaid can pay for health AI that makes sure that equity is considered in their purchasing or reimbursement technologies, the NIH [National Institute of Health] can fund more research in making algorithms more equitable and build standards for these algorithms that could then be used by the FDA [Food and Drug Administration] as they're trying to figure out what health equity means and how they're regulated within their current authorities," they continue.

Among others, the paper lists six primary existing government agencies that could help regulate health AI, including: the FDA, the Federal Trade Commission (FTC), the recently established Advanced Research Projects Agency for Health, the Agency for Healthcare Research and Quality, the Centers for Medicare and Medicaid, the Department of Health and Human Services, and the Office of Civil Rights (OCR).

But Wickerson says that more must be done. The most challenging part to writing the paper, in Wickerson's view, was "imagining what we don't have yet."

Rather than solely relying on existing [regulatory bodies](#), the paper also proposes creating an independent auditing authority, similar to the NTSB, that allows for a safety audit for malfunctioning health AI systems.

"I think that's the current question for tech governance—we haven't really had an entity that's been assessing the impact of technology since

the '90s," Wickerson adds. "There used to be an Office of Technology Assessment ... before the digital era even started, this office existed and then the federal government allowed it to sunset."

Zach Harned, co-author and recent graduate of Stanford Law School, believes a primary challenge in emerging technology is having technological development outpace regulation. "However, the importance of AI technology and the potential benefits and risks it poses, especially in the [health-care](#) arena, has led to a flurry of regulatory efforts," Harned says.

"The FDA is clearly the primary player here, and they've consistently issued guidances and white papers attempting to illustrate their evolving position on AI; however, privacy will be another important area to watch, with enforcement from OCR on the HIPAA [Health Insurance Portability and Accountability Act] side and the FTC enforcing privacy violations for non-HIPAA covered entities."

Harned notes that the area is evolving fast, including developments such as the recent White House Executive Order 14110 on the safe and trustworthy development of AI, as well as regulatory activity in the European Union (EU), including the capstone EU AI Act that is nearing finalization. "It's certainly an exciting time to see this important technology get developed and regulated to ensure safety while also not stifling innovation," he says.

In addition to regulatory activities, the paper suggests other opportunities to create incentives for safer health AI tools such as a pay-for-performance program, in which insurance companies reward hospitals for good performance (though researchers recognize that this approach would require additional oversight to be equitable).

So just how long do researchers think it would take to create a working

regulatory system for health AI? According to the paper, "the NTSB and FAA system, where investigations and enforcement are in two different bodies, was created by Congress over decades."

Bondi-Kelly hopes that the paper is a piece to the puzzle of AI regulation. In her mind, "the dream scenario would be that all of us read the paper and are super inspired and able to apply some of the helpful lessons from aviation to help AI to prevent some of the potential harm that might come about."

In addition to Ghassemi, Shah, Bondi-Kelly, and Sanneman, MIT co-authors on the work include Senior Research Scientist Leo Anthony Celi and former postdocs Thomas Hartvigsen and Swami Sankaranarayanan.

More information: Elizabeth Bondi-Kelly et al, Taking Off with AI: Lessons from Aviation for Healthcare, *Equity and Access in Algorithms, Mechanisms, and Optimization* (2023). [DOI: 10.1145/3617694.3623224](https://doi.org/10.1145/3617694.3623224)

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: Stratospheric safety standards: How aviation could steer regulation of AI in health (2024, January 17) retrieved 29 May 2024 from <https://techxplore.com/news/2024-01-stratospheric-safety-standards-aviation-ai.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.