The Harvard Gazette

HEALTH

AI is speeding into healthcare. Who should regulate it?

Medical ethicist details need to balance thoughtful limits while avoiding unnecessary hurdles as industry groups issue guidelines

Alvin Powell | Harvard Staff Writer

January 12, 2026 ▪ 8 min read



I. Glenn Cohen.
File photo by Niles Singer/Harvard Staff Photographer

AI is moving quickly into healthcare, bringing potential benefits but also possible pitfalls such as bias that drives unequal care and burnout of physicians and other healthcare workers. It remains undecided how it should be regulated in the U.S.

In September, the hospital-accrediting **Joint Commission** and the **Coalition for Health AI** issued **recommendations** for implementing artificial intelligence in medical care, with the burden for compliance falling largely on individual facilities.

I. Glenn Cohen, faculty director of **Harvard Law School**'s **Petrie-Flom Center for Health Law, Biotechnology, and Bioethics**, and colleagues suggested in the **Journal of the American Medical Association** that the guidelines are a good start, but changes to ease likely regulatory and financial burdens — particularly on small hospital systems — are needed.

In this edited conversation, Cohen, the James A. Attwood and Leslie Williams Professor of Law, discussed the difficulty of balancing thoughtful regulation with avoiding unnecessary roadblocks to game-changing innovation amid rapid adoption.

Is it clear that AI in healthcare needs regulation?

Whenever medical AI handles anything with medium to high risk, you want regulation: internal self-regulation or external governmental regulation. It's mostly been internal thus far, and there are differences in how each hospital system validates, reviews, and monitors healthcare AI.

When done on a hospital-by-hospital basis like this, costs to do this kind of evaluation and monitoring can be significant, which means some hospitals can do this, and some can't. By contrast, top-down regulation is slower — maybe too slow for some forms of progress in this space.

There's also a complicated mix of AI products going into hospitals. Some may assist with things like internal purchasing and review, but many more are clinical or clinically adjacent.

Some medical AI products interface directly with consumers, such as chatbots that people might be using for their mental health. For that, we don't even have internal hospital review, and the need for regulation is much clearer.

With technology moving so fast, is speed important even in regulation?

This is an innovation ecosystem that has a lot of startup energy, which is great. But you're talking about something that can scale extremely quickly, without a lot of internal review.

Whenever you enter what I call a "race dynamic," there is a risk that ethics is left behind pretty quickly. Whether the race is to be the first to develop something, a race for a startup against money running out, or a national race between countries trying to develop artificial intelligence, the pressures of time and urgency make it easier to overlook ethical issues.

The vast majority of medical AI is never reviewed by a federal regulator — and probably no state regulator. We want to have standards for healthcare AI and an incentive to adopt standards.

But putting everything through the rigorous FDA process for drugs or even the one for medical devices would in many cases be prohibitively expensive and prohibitively slow for those enamored with the rate of development in Silicon Valley.

On the flip side, if they perform badly, many of these technologies are a much greater risk to the general populace than the average device on the market.

If you take an aspirin or a statin, there are differences in how they work in different people, but to a large extent we can characterize those differences ahead of time. When medical AI is reading an X-ray or doing something in the mental health space, how it's implemented is key to its performance.

You might get very different results in different hospital systems, based on resources, staffing, training, and the experience and age of people using them, so one has to study implementation very carefully. This would create an unusual challenge for an agency like FDA — which often says it does not regulate the practice of medicine — because where the approval of an AI system stops and the practice of medicine begins is complicated.

Your study examines a regulatory system suggested by the Joint Commission, a hospital accreditor, and the Coalition for Health AI. Would an accreditor naturally be something that hospitals would — or would have to — pay attention to?

Exactly. In almost every state, in order to be able to bill Medicare and Medicaid you need to be accredited by the Joint Commission. This is a huge part of almost every hospital's business.

There is a robust process to qualify for accreditation, and every so often you are re-evaluated. It's serious business.

The Joint Commission hasn't yet said that these AI rules are going to be part of our next accreditation, but these guidelines are a sign that they may be going in that direction.

> “I speak about legal and ethical issues in this space, but I’m an optimist about this. I think that, in 10 years, the world will be significantly better off because of medical artificial intelligence.”

Do you find some of the recommendations wanting?

Some are more demanding than I expected, but I actually think they’re pretty good.

Requiring that — when appropriate — patients should be notified when AI directly impacts their care and that — when relevant — consent to use an AI agent should be obtained, is a strong position to take.

A lot of scholars and other organizations don’t take the position that medical AI should always be disclosed when it directly impacts care, let alone that informed consent should always be sought.

The guidelines also require ongoing quality monitoring and continual testing, validation, and monitoring of AI performance.

Monitoring frequency would scale to risk levels in patient care. These are good things to do, but difficult and expensive. You’ll have to assemble multidisciplinary AI committees and constantly measure for accuracy, errors, adverse events, equity, and bias across populations.

If taken seriously, it will probably be infeasible for many hospital systems in the U.S. They will have to make a threshold decision whether they’re going to be AI adopters.

You point out in your JAMA article that most hospitals in the U.S. are small community hospitals, and that resources are a major issue.

I am told by people in major hospital systems that already do this that to properly vet a complex new algorithm and its implementation can cost $300,000 to half a million dollars. That's simply out of reach for many hospital systems.

There are actually going to be things in the implementation that are specific to each hospital, but there are also going to be things that might be valuable to know that are common for many hospital systems. The idea that we'd do the evaluation repeatedly, in multiple places, and not share what's learned seems like a real waste.

If the answer is, "If you can't play in the big leagues, you shouldn't step up to bat," that creates a have/have-not distribution in terms of healthcare access. We already have that as to healthcare generally in this country, but this would further that dynamic at the hospital level.

Your access to AI that helps medical care would be determined by whether you're in networks of large academic medical centers that proliferate in places like Boston or San Francisco, rather than other parts of the country that don't have that kind of medical infrastructure.

The goal, ideally, would be more centralization and more sharing of information, but these recommendations put a lot of the onus on individual hospitals.

Doesn't a system where some hospitals can't participate negate the potential benefits from this latest generation of AI, which can assist places that are resource-poor by providing expertise that might be missing or hard to find?

It would be a shame if you've got a great AI that's helping people and might do the most benefit in lower-resource settings, and yet those settings are unable to meet the regulatory requirements in order to implement.

It would also be a sad reality, as an ethical matter, if it turns out that we're training these models on data from patients across the country, and many of those patients will never get the benefit of these models.

If the answer is that the vetting and monitoring of medical AI should be done by a larger entity, is that the government?

The Biden administration's idea was to have "assurance labs" — private-sector organizations that in partnership with the government could vet the algorithms under agreed-upon standards such that healthcare organizations could rely on them.

The Trump administration agrees on the problem but has signaled that they don't like the approach. They have yet to fully indicate what their vision is.

It sounds like a complex landscape, as well as a fast-moving one.

Complex, but also challenging and interesting.

I speak about legal and ethical issues in this space, but I'm an optimist about this. I think that, in 10 years, the world will be significantly better off because of medical artificial intelligence.

The diffusion of those technologies to less-resourced settings is very exciting, but only if we align the incentives appropriately. That doesn't happen by accident, and it is important that these distributional concerns be part of any attempt to legislate in the area.



HARVARD
UNIVERSITY