The Robot Doctor Will See You Now (Feb 2, 2025)

By Pranav Rajpurkar and Eric J. Topol

Dr. Rajpurkar researches the application of A.I. in clinical settings. Dr. Topol is a cardiologist.

The rapid rise in artificial intelligence has created intense discussions in many industries over what kind of role these tools can and should play — and health care has been no exception. The medical community largely anticipated that combining the abilities of doctors and A.I. would be the best of both worlds, leading to more accurate diagnoses and more efficient care.

That assumption might prove to be incorrect. A growing body of research suggests that A.I. is outperforming doctors, even when they use it as a tool.

A recent M.I.T.-Harvard study, of which one of us, Dr. Rajpurkar, is an author, examined how radiologists diagnose potential diseases from chest X-rays. The study found that when radiologists were shown A.I. predictions about the likelihood of disease, they often undervalued the A.I. input compared to their own judgment. The doctors stuck to their initial impressions even when the A.I. was correct, which led them to make less accurate diagnoses. Another trial yielded a similar result: When A.I. worked independently to diagnose patients, it achieved 92 percent accuracy, while physicians using A.I. assistance were only 76 percent accurate — barely better than the 74 percent they achieved without A.I.

This research is early and may evolve. But the findings more broadly indicate that right now, simply giving physicians A.I. tools and expecting automatic improvements doesn't work. Physicians aren't completely comfortable with A.I. and still doubt its utility, even if it could demonstrably improve patient care.

But A.I. will forge ahead, and the best thing for the medicine to do

is to find a role for it that doctors can trust. The solution, we believe, is a deliberate division of labor. Instead of forcing both human doctors and A.I. to review every case side by side and trying to turn A.I. into a kind of shadow physician, a more effective approach is to let A.I. operate independently on suitable tasks so that physicians can focus their expertise where it matters most.

Sign up for the Opinion Today newsletter Get expert analysis of the news and a guide to the big ideas shaping the world every weekday morning.

What might this division of labor look like? Research points to three distinct approaches. In the first model, physicians start by interviewing patients and conducting physical examinations to gather medical information. A Harvard-Stanford study that Dr. Rajpurkar helped write demonstrates why this sequence matters — when A.I. systems attempted to gather patient information through direct interviews, their diagnostic accuracy plummeted — in one case from 82 percent to 63 percent. The study revealed that A.I. still struggles with guiding natural conversations and knowing which follow-up questions will yield crucial diagnostic information. By having doctors gather this clinical data first, A.I. can then apply pattern recognition to analyze that information and suggest potential diagnoses.

In another approach, A.I. begins with analyzing medical data and suggesting possible diagnoses and treatment plans. A.I. seems to have a natural penchant for such tasks: A 2024 study showed that OpenAI's latest models perform well at complex critical thinking tasks like generating diagnoses and managing health conditions when tested on case studies, medical literature and patient scenarios. The physician's role is to then apply his clinical judgment to turn A.I.'s suggestions into a treatment plan, adjusting the recommendations based on a patient's physical limitations, insurance coverage and health care resources.

The most radical model might be complete separation: having A.I.

handle certain routine cases independently (like normal chest X-rays or low-risk mammograms), while doctors focus on more complex disorders or rare conditions with atypical features.

Early evidence suggests this approach can work well in specific contexts. A Danish study published last year found that an A.I. system could reliably identify about half of all normal chest X-rays, freeing up radiologists to devote more time to studying images that were deemed suspicious. In a landmark Swedish trial involving mammograms for more than 80,000 women, half the scans were assessed by two radiologists, as is usual. The other half were evaluated by A.I.-supported screening first, followed by additional review by one radiologist (and in rarer instances where the A.I. determined an elevated risk, by two radiologists). The A.I.-assisted approach led to the identification of 20 percent more breast cancers while reducing the overall radiologist workload almost in half.

This might be the clearest path to dealing with the shortage of health care workers hurting medicine. This model is particularly promising for underserved areas, where A.I. systems could provide initial screening and triage, so limited specialist resources can be redirected to more pressing issues.

All these approaches raise questions about liability, regulation and the need for ongoing clinician education. Medical training will need to adapt to help doctors understand not just how to use A.I., but when to rely on it and when to trust their own judgment. Perhaps most important, we still lack definitive proof that these approaches, tested in research studies or pilot programs, will achieve the same success in the messy realities of everyday care.

But the promise for patients is obvious: fewer bottlenecks, shorter waits and potentially better outcomes. For doctors, there's potential for A.I. to alleviate the routine burdens so that health care might become more accurate, efficient and — paradoxically

— more human.

Pranav Rajpurkar is an assistant professor at Harvard Medical School and a founder of the company a2z Radiology AI. Eric J. Topol is a professor and executive vice president at Scripps Research in La Jolla, Calif.