THE ARTIFICIAL INTELLIGENCE ISSUE

THE FUTURE OF EVERYTHING

A LOOK AHEAD FROM THE WALL STREET JOURNAL.

The Brewing Solar Storm

Scientists are building early-warning systems to help protect Earth R4

AI's Mind-Body Problem

A science fiction writer on technology's future forms R2

A New Eye On Alzheimer's

Scans aim to detect disease years before symptoms R9

ILLUSTRATION BY JACK

MEDICAL DEVICES GET SMARTER. CAN THE FDA KEEP UP?

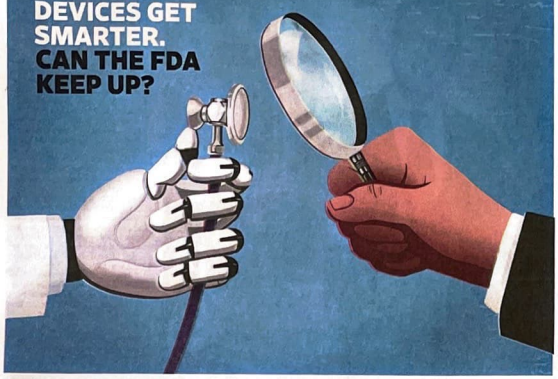

Use of AI means devices and apps can learn and change, making regulators' job harder **By Ryan Tracy**

The use of artificial intelligence in medical devices is exploding, paving the way for potentially lifesaving advances—and daunting challenges for regulators responsible for keeping new products safe.

If tech companies have their way, anxious patients may one day opt to talk to a chatbot instead of a human therapist. A patient checking into the hospital would be assessed with an AI-driven risk calculator, and an automated assistant would help do tests. *Such AI-powered apps, when used for diagnosis or treatment, would qualify as* medical devices under the oversight of the Food and Drug Administration—meaning developers could need the agency's permission to market them.

The FDA is one of many U.S. agencies wrestling to find the right policies as the AI revolution picks up speed. From 2020 through late 2022, the FDA approved more than 300 devices with AI features—more than during the entire previous decade. "We've got a speeding train coming in a tunnel and there's a challenge about to hit," says Amy Abernethy, who served as the FDA's principal deputy commissioner of food and drugs from 2019 to 2021 and is now chief medical officer at Verily Life Sciences, a unit of Alphabet that makes AI-enabled medical devices.

For decades, the agency looked at medical devices the same way it looks at drugs: as static compounds. When the FDA approves a device, the manufacturer can sell that version. It needs the regulator's signoff before upgrading to a new version. But AI-enabled devices often use algorithms designed to be updated rapidly, or even learn on their own. The agency is grappling with how to deal with this fast-moving technology while ensuring the devices stay safe.

"One of the key benefits of an AI-powered product is that it can be improved over time," says Tim Sweeney, CEO of Inflammatix, a developer of blood tests designed to predict the presence, type, and severity of an infection. If the company learned that a particular pattern of the body's immune response strongly indicates the onset of sepsis, for instance, it would want to retrain its algorithms to account for that. "If you have a lot of extra data, you should be improving your results," Sweeney says.

Under the FDA's traditional method of oversight, however, companies like Inflammatix would likely have to get additional permission before changing their algorithms.

The agency is beginning to point developers down an alternative path, in April offering formal guidance on how they can submit more flexible *plans for devices that use AI. A manufacturer can file a "Predetermined* Change Control Plan" that outlines expected alterations. FDA staff—including lawyers, doctors, and tech experts—review the plans and the scope of the expected changes. Once the device is approved, the company can alter the product's programming without the FDA's blessing, as long as the changes were part of the plan.

The goal is "to ensure that the market has some flexibility to continue to push the boundaries of the innovation that they are doing, while still adhering to the guardrails that we've established when we approve a product," says Troy Tazbaz, director of the FDA's Digital Health Center of Excellence.

Many device makers have cheered the FDA direction, though they have asked the agency for more clarity on the extent of changes it would allow. "It's going to be challenging to, up front, predetermine what are the possible changes that we would want to do," says Lara Jehi, who has helped develop AI-driven patient assessment tools as chief research information officer for the Cleveland Clinic Health System. "The only thing we know is predictable in science is that it is unpredictable."

NeuroPace, a maker of brain implants that treat severe epilepsy, would consider a "change control plan" as part of a new AI-enabled version of its device, says Martha Morrell, the company's chief medical officer.

The device now delivers electrical currents to stop a seizure, and those currents are only altered by a physician after reviewing the patient's data. In the future, Morrell says, the device could be given the power to interpret the patient's brain waves and decide what current to deliver in real time, subject to predetermined safety limits.

'We've got a speeding train coming in a tunnel and there's a challenge about to hit.'
Amy Abernethy, Verily Life Sciences

Morrell says that, as a "type-A doctor," even she feels uncomfortable ceding that kind of power to a device. But she also believes patients will benefit and that the FDA and the company need to figure out the standard necessary to be confident that the device will be safe.

After the latest FDA guidance on AI devices, she predicts, "there will be another one right behind it, and another one right behind that as well, as we scramble to keep right on top of this enormous change."

Critics see the FDA's new direction as potentially introducing risks, pointing out that there have been cases where medical-device software mistakes posed grave dangers to patients. AI systems add an additional challenge, they say, because the humans in charge of them don't always fully understand how they work and potential problems, such as racial bias, can be difficult to measure. "AI brings with it very different kinds of harms—algorithmic harms that are difficult to understand and explain," says Mason Marks, a health-law professor at Florida State University. By allowing changes to complex algorithms without going through regulatory scrutiny, he says, "the FDA is giving the manufacturers a lot of leeway."

An FDA spokesman said the agency's evaluation of devices addresses potential bias.

Those concerns could, in the future, be addressed by a new policy some FDA officials have discussed: requiring real-time monitoring of AI devices after they go on the market. That would look different from the traditional approach, where the maker of an FDA-approved device must simply report "adverse events" that could affect patients. New monitoring requirements would aim to ensure manufacturers are closely watching their algorithms' performance. But adopting them would likely require legal authority the FDA doesn't have. Industry groups might lobby against the agency gaining that authority, if they view the requirements as overly intrusive.

These debates will only increase in urgency as more developers incorporate generative AI systems that can produce humanlike outputs of text, photos, videos and other media.

Woebot Health offers a bot that it hopes psychiatrists, doctors, or insurers will recommend to help patients work on mental-health issues. The bot offers encouragement and responds 24/7. As a safety measure, Woebot says, its responses are limited to a selection preapproved by humans. The company is testing the use of generative AI model to help understand what users say, but is wary of allowing open conversations.

Woebot markets its chatbot as a mental-health support app, avoiding using certain terms, such as "treatment," that could trigger the need for an FDA review. The agency has authority to regulate any device intended to be used for the diagnosis, treatment or prevention of disease—a broad remit that applies to many software functions. It doesn't enforce its rules in what it views as low-risk situations, such as wellness apps.

Woebot says it wants to seek the FDA's stamp of approval. But it also wants to be able to market its product without a prescription, reaching a much broader user base. FDA-approved "digital therapeutics" for mental health generally aren't cleared for sales directly to consumers. "How can we make sure these tools are clinically validated [but do] not necessarily require a prescription?" says Robbert Zuster zeel, Woebot's vice president of regulatory science and strategy. "There's this middle option at some point the FDA needs to look at."

That points to another potential effect of the FDA's policies: The agency's regulatory boundaries end up shaping product development.

"It creates a funny set of incentives," says Ariel Dora Stern, a professor at Harvard Business School who has studied the agency. "I've talked to startups that are basically going after smaller problems than their algorithms can go after, simply because regulation is something that they don't want to go near with a 10-foot pole."

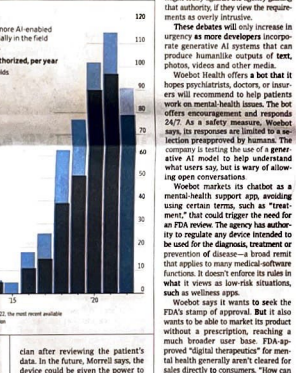

AI Psychotherapist

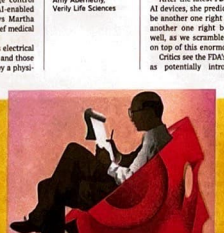

Financial firms relying on AI for prediction and decision-making will need people to divine the drivers of a model's thinking.

Unlike conventional software, the logic behind the output of applications such as OpenAI's ChatGPT is typically opaque. That may be fine when they're used to generate things like recipes and poems, says Dinesh Nirmal, senior vice president, products, IBM Software, but not if they're relied upon for things such as assigning credit scores, optimizing investment portfolios and predicting liquidity balances.

Business or enterprise AI, which serves firms and organizations, is all about "explainability," says Nirmal. Customers will want to know why their loan application was rejected. Bank regulators will require some decisions to be explained.

AI psychotherapists will evaluate a model's upbringing, by scrutinizing its training data for errors and sources of bias. They may put AI models on the couch, by probing them with test questions. Companies such as IBM, Google and Microsoft are racing to release tools that quantify and chart an AI's thought processes, but like Rorschach tests they require people to interpret their outputs.

Understanding an AI's reasoning will only be half the job, says Alexey Surkov, partner and global head of model risk management at Deloitte & Touche. The other half will be signing off on a model's mental fitness for the task at hand. "No matter how sophisticated the models and systems get," says Surkov, "we as humans are ultimately responsible for the outcomes of the use of those systems."

"Psychotherapist" might be a stretch for a job title in some financial firms. Surkov suggests AI Risk Manager or Controller as alternatives. Join

—Bob Henderson

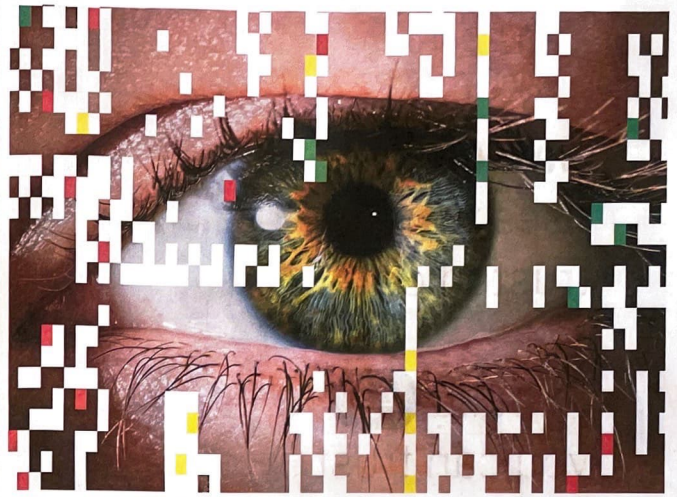

THE EYE AS A WINDOW TO ALZHEIMER'S

AI tools could diagnose the disease with visual scans, years before symptoms appear
By Vipal Monga

Getting tested for Alzheimer's disease could one day be as easy as checking your eyesight.

RetiSpec has developed an artificial-intelligence algorithm that it says can analyze results from an eye scanner and detect signs of Alzheimer's 20 years before symptoms develop. The tool is part of broader work by startups and researchers to harness AI to unlock the mysteries of a disease that afflicts more than seven million Americans.

For years, people have studied individual hallmarks of Alzheimer's, including brain inflammation and neurodegeneration, but the exact causes of the disease remain elusive. AI, researchers say, could open a new era in the diagnosis of a neurological disease that remains difficult to identify, let alone treat.

"There still remains a huge amount we fundamentally don't understand about the brain and how it works," said Elias Shaked, co-founder of Torrentio-based RetiSpec. "The power in AI is that it can help connect the dots."

Another company, Sacramento, Calif.-based Neurovision, aims to use machine learning to develop retinal scans and blood tests to identify people at risk of developing Alzheimer's and other forms of dementia. The company's AI model analyzes eye scans for anomalies, like the buildup of certain proteins or blood vessels with a twisted shape, that are associated with Alzheimer's, said Steven Verdooner, Neurovision's co-founder.

It can be difficult for people to discern such signs in the scans. Many scans have dark areas, and plaque deposits can be very small. The human eye can't distinguish them very well, Verdooner said. "The algorithm does a better job," he said.

At the University of Arizona College of Medicine in Tucson, neurology associate professor Rui Chang built an AI model that aims to identify genetic triggers that are linked to Alzheimer's. The traditional approach researchers follow is painstakingly slow, Chang said.

"It's like looking at the forest one tree at a time," he said.

AI can absorb the whole forest of information at once and find patterns people can't. The model took two months to identify 6,000 gene targets that, if knocked down or repressed, might change how Alzheimer's develops, Chang said the tool has cut a decade off his research.

Chang founded a company, Path Biotech, which will start clinical trials next year based on his AI research.

Alzheimer's was the sixth-leading cause of death in the U.S. in 2021, not including Covid-19.

The Food and Drug Administration in July approved a drug, Leqembi, that removes amyloid, a sticky plaque that gathers in the brains of Alzheimer's patients. But current techniques for identifying the disease are expensive and difficult.

People with symptoms can get a spinal tap or a PET scan to see if they have high levels of amyloid and tangled strands of the protein tau, which is also commonly found in Alzheimer's patients. The scans are very accurate, which makes them the gold standard of diagnosis, said Catherine Bornbaum, RetiSpec's chief business officer. Compared with autopsy results—still the only way to tell for certain whether a patient had Alzheimer's when they died—the PET scan's diagnostic success rate is close to 90%.

Left, unenhanced retinal scan. Right, a scan with Neurovision's AI shows white specks indicating the presence of the amyloid protein

80%
accuracy of RetiSpec's AI in detecting optical signatures that correspond with amyloid in the brain, a study found

But the machines aren't widely available, the scans are expensive, and getting a diagnosis can take weeks. Insurers don't routinely cover the scans and they can cost around $6,000.

Artificial intelligence technologies could speed up the process of getting a diagnosis and make it cheaper. RetiSpec's AI reads scans from a camera that can be attached to machines already available in most optometrists' offices, for example. The camera measures a wider range of the spectrum than the human eye can see, which allows the AI to detect unique optical signatures that correspond with the presence of amyloid in the brain. The model, which delivers results instantaneously, was 80% accurate in detecting such signatures in a recent study of 271 patients.

AI tools in medical research can perform well in clinical testing but break down in messier real-life situations, said Matt Leming, a postdoctoral researcher at Massachusetts General Hospital.

"Biotech AI models are finicky," he said.

AI learns better from huge amounts of data, Leming said. AI models like ChatGPT are good at analyzing and mimicking writing, for example, because they learn from text gathered across the internet.

Medical data is comparatively scarce and proprietary. That means AI in biotech has a more limited sample to learn from and its results can be easily thrown off by wider variation in cases it encounters in a clinic compared with more controlled laboratory settings.

"When it comes to AI fundamentally changing the way we do medicine, I don't think it is going to happen," Leming said.

Chang at the University of Arizona said he has tried to overcome this problem by using mathematical models that minimize errors and improve prediction accuracy. RetiSpec said the company has taken samples from 14 research partners, from whom it gathers samples from racially and socioeconomically diverse communities. Neurovision said it took samples from diverse data sets and tested them against others to minimize errors.

"Some of the most important work we've done is to be sure that AI doesn't suffer from garbage in and garbage out," RetiSpec's Shaked said.

Reskiller

As AI becomes capable of taking on more work that is now done by humans, people will need to more aggressively upgrade their skills to stay productive and employable. "Reskillers," a new type of teacher, will help people stay one step ahead of the machines.

Companies will put growing value on specialists who can guide such critical human development. "Teachers had it bad under the industrial revolution. Look at what they are paid," says Stephen Messer, co-founder and chairperson of Collective[i], which has developed a foundation model that produces insights around revenue forecasting and growth. "Now, I think teachers are about to go through a revolution because of AI."

Reskillers will need to understand the talents that organizations require as technology marches ahead. "This puts an onus on employees and companies to stay relevant," says Keith Peiris, co-founder and chief executive of Torne, a startup with a generative AI-native storytelling and presentation platform. "In the 'old world,' pre generative AI, maybe you needed 100 people to build a company. With AI, maybe you could build that company with 30 people."

New career-development arcs are already taking shape because of generative AI, according to Peiris, who is trained in nanobiotechnology engineering and is a veteran of Facebook and Instagram. "Sales professionals are learning web development and copywriting. Marketers are becoming more steeped in graphic design," he says. "HR professionals are taking on "legal work" and becoming paralegals by using AI legal tools."

—Steven Rosenbush

FROM TOP: PHOTO ILLUSTRATION BY JOEL KIMMEL; NEUROVISION CO.; IRINA CHAPOVA; ISTOCK (2019)