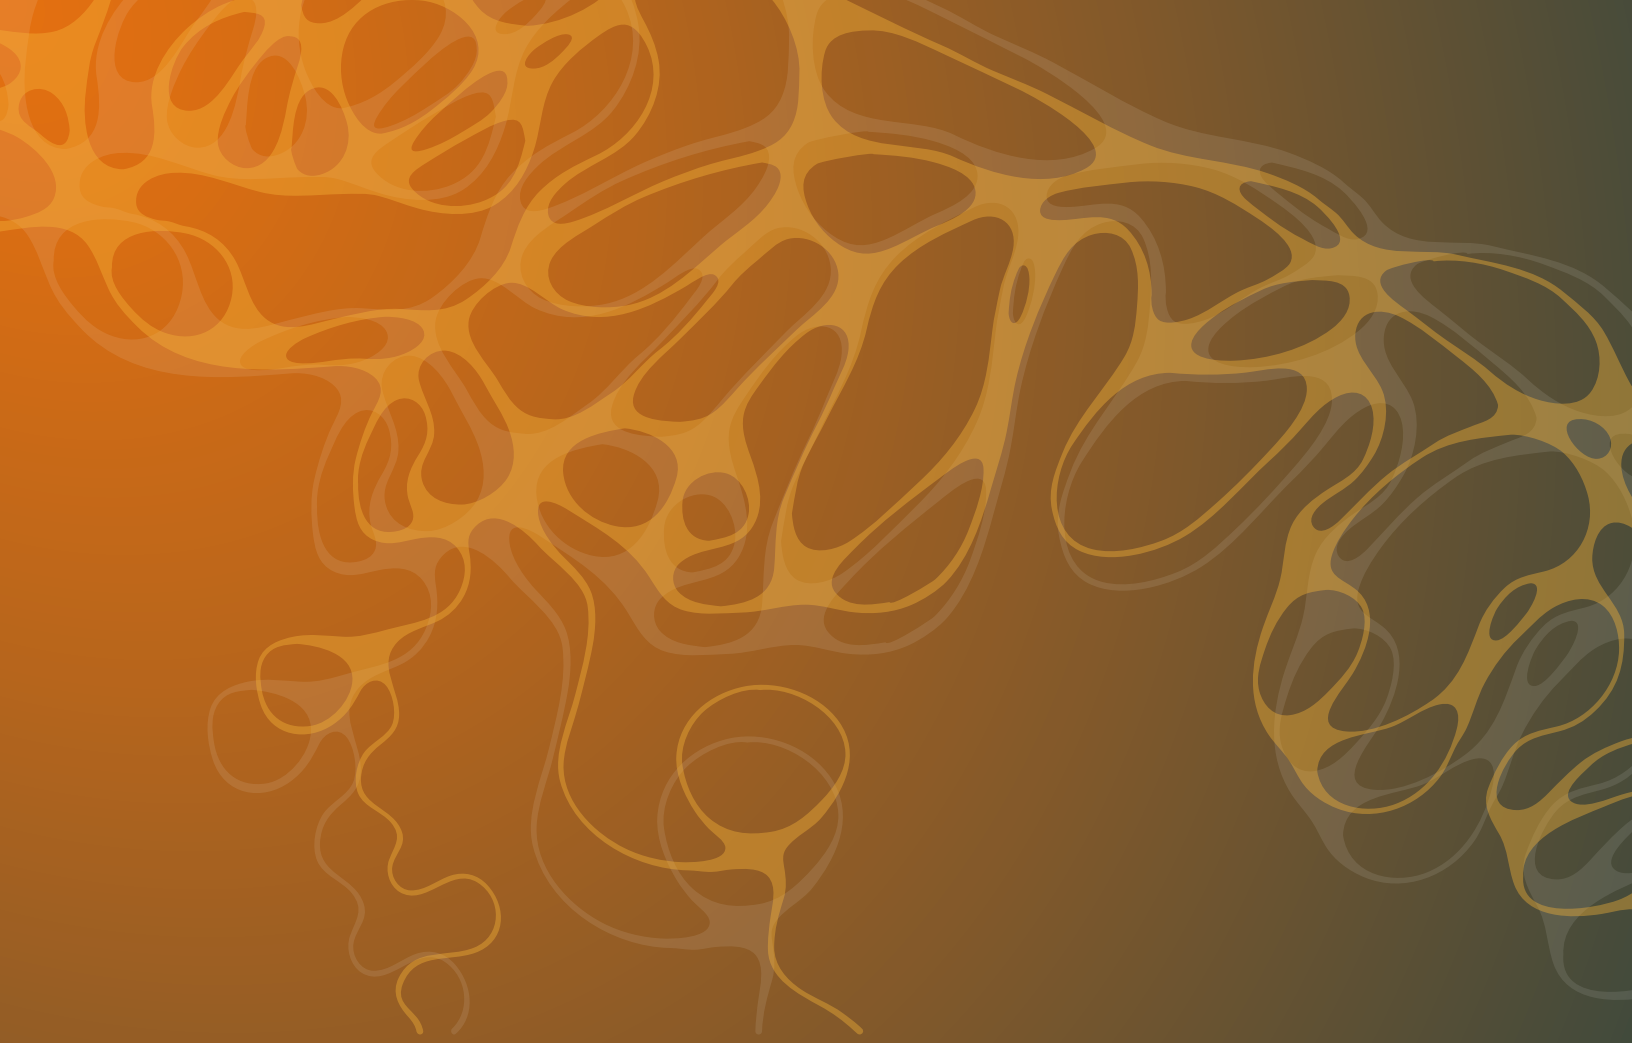

STAT E-BOOK

AI applications in clinical research

SPONSORED BY IMOHealth

STAT

TABLE OF CONTENTS

STAT

In medical parlance, "stat" means important and urgent, and that's what we're all about — quickly and smartly delivering good stories. We take you inside science labs and hospitals, biotech boardrooms, and political backrooms. We dissect crucial discoveries. We examine controversies and puncture hype. We hold individuals and institutions accountable. We introduce you to the power brokers and personalities who are driving a revolution in human health. These are the stories that matter to us all.

BOSTON • WASHINGTON • NEW YORK • LONDON

SAN FRANCISCO • LOS ANGELES • CLEVELAND • CHICAGO

Our team includes talented writers, editors, and producers capable of the kind of explanatory journalism that complicated science issues sometimes demand. And even if you don't work in science, have never stepped foot in a hospital, or hated high school biology, we've got something for you. The world of health, science, and medicine is booming and yielding fascinating stories. We explore how they affect us all. And, with our eBook series, we regularly do deep dives into timely topics to get you the inside scoop you need.

Introduction

Artificial intelligence (AI) is no longer just a buzzword in clinical research; it is a transformative force that shapes how we understand, diagnose, and treat disease and practice medicine. The convergence of machine learning, vast datasets, and ever-evolving algorithms is driving innovations that once seemed impossible. This ebook dives into several of the real-world applications of AI in clinical research, as reported by STAT. These (and many other) challenges have been reported on in STAT, and are highlighted in this collection — from the absence of medically tailored food delivery to Medicare recipients to AI-based denial of insurance coverage.

We begin with cutting-edge work where researchers have employed machine learning to translate brain signals from a paralyzed patient into text. By implanting electrodes in the patient's brain, UCSF scientists successfully interpreted neural signals related to speech production, achieving a 75% accuracy rate in generating words and sentences.

Next, we consider AI's role in enhancing the human side of medicine. Peter Lee, who oversees Microsoft's health care initiatives and spoke at the 2024 STAT Breakthrough West Summit, views generative AI as a transformative tool with the potential to revolutionize medicine. While AI models like GPT-4 are not perfect, he says, they have the ability to enhance patient care by improving tasks like after-visit summaries, often adding a personal touch that resonates with patients.

Pulitzer-nominated National Technology Correspondent Casey Ross, STAT's in-house expert on all things AI, reports on the rush to integrate generative AI into operations, with major players like Microsoft and Nvidia leading the charge. Despite the excitement, experts warn that the hype around AI's potential in health care could lead to unrealistic expectations, similar to past overpromises like IBM's Watson.

The power of AI extends far beyond urban centers and academic hospitals. In a First Opinion piece, Bill Gassen emphasizes that rural health care providers are pioneering the use of AI — using it to streamline administrative tasks, reduce cognitive burdens, and enhance clinical decision-making.

The financial influx into AI-driven drug discovery raises critical questions: will billions of dollars in funding lead to a proportional increase in new treatments? The authors of this piece argue that while AI is revolutionizing drug discovery, focusing solely on this aspect without reimagining the entire pharmaceutical R&D system risks bottlenecks and inefficiencies.

As AI becomes increasingly integrated into health care, the need for thoughtful regulation grows. Regulators, industry bodies, and lawmakers are struggling to agree on frameworks to ensure safety, with some advocating for traditional regulatory approaches while others suggest self-regulation through insurance and reparations for harm caused by AI. Experts argue that, like other industries, health care AI should be regulated systematically, drawing parallels to how safety standards in aviation and automotive industries have historically been implemented through regulation.

And in counterpoint, Garry Tan — CEO of Y Combinator — and other tech leaders are advocating for flexible AI regulations that don't stifle innovation, particularly for smaller startups.

Through these articles, this ebook provides a deep dive into AI's transformative role in clinical research. Whether you are a researcher, clinician, or industry leader, the insights offered here will deepen your understanding of how AI is redefining the future of medicine.

For more, visit statnews.com/topic/artificial-intelligence/.

– Jesse McQuarters, Editor, STAT Brand Studio

IMO Health

Transforming data. Transforming healthcare.

IMO Health is a clinical data intelligence business at the heart of a digital revolution in healthcare. Combining robust medical terminology, extensive domain knowledge, and AI, we expertly structure and operationalize clinical data to generate sharper insights and inform more intelligent decision-making.

Learn more at imohealth.com

IN THE LAB

STAT

Researchers use machine learning to translate brain signals from a paralyzed patient into text

By Claudia Lopez Lloreda | **JULY 14, 2021**

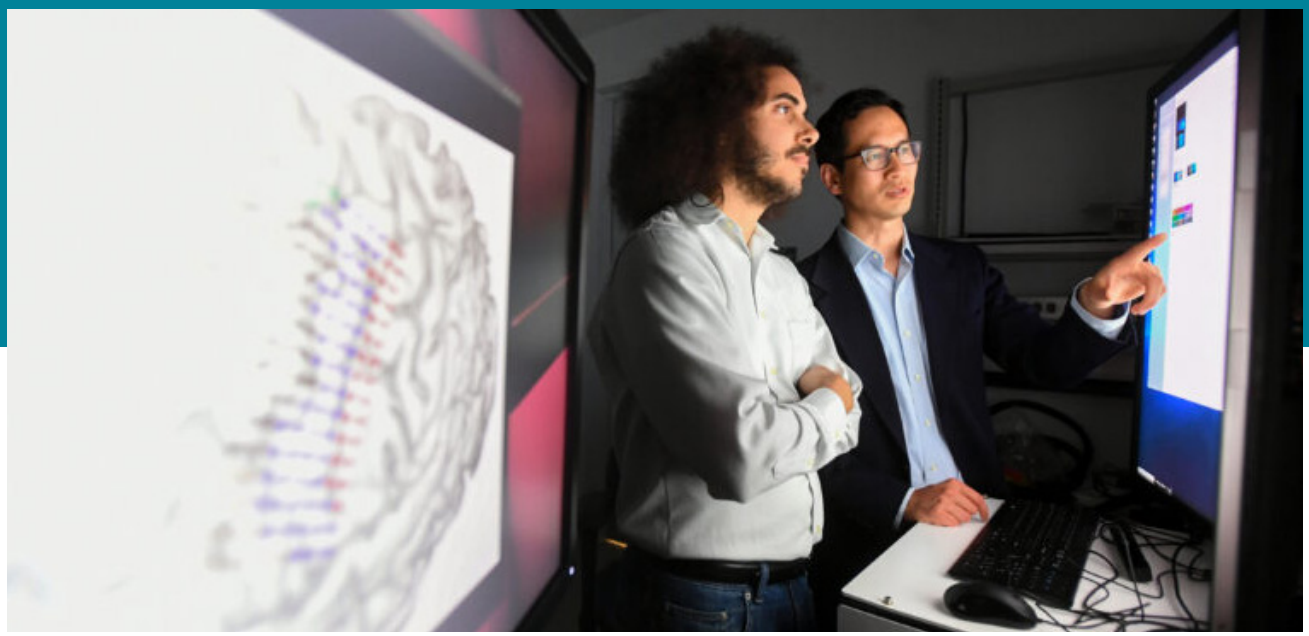

David Moses (left) and his adviser at UCSF's Mission Bay campus in San Francisco.
COURTESY NOAH BERGER

Assistive technologies such as handheld tablets and eye-tracking devices are increasingly helping give voice to individuals with paralysis and speech impediments who otherwise would not be able to communicate. Now, researchers are directly harnessing electrical brain activity to help these individuals.

In a study published Wednesday in the New England Journal of Medicine, researchers at the University of California, San Francisco, describe an approach that combines a brain-computer interface and machine learning models that allowed them to generate text from the electrical brain activity of a patient paralyzed because of a stroke.

Other brain-computer interfaces, which transform brain signals into commands, have used neural activity while individuals attempted handwriting movements to produce letters. In a departure from previous work, the new study taps into the speech production areas of the brain to generate entire words and sentences that show up on a screen.

This may be a more direct and effective way of producing speech and helping patients communicate than using a computer to spell out letters one by one, said David Moses, a UCSF postdoctoral researcher and first author of the paper.

The study was conducted in a single 36-year-old patient with anarthria, a condition that renders people unable to articulate words because they lose control of muscles tied to speech, including in the larynx, lips, and tongue. The anarthria was brought on by a stroke more than 15 years ago that paralyzed the man.

The researchers implanted an array of electrodes in the patient's brain, in the area that controls the vocal tracts, known as the sensorimotor cortex. They measured the electrical activity in the patient's brain while he was trying to say a word and used a machine learning algorithm to then match brain signals with specific words. With this code, the scientists prompted the patient with sentences and asked him to read

them, as though he were tying to say them out loud. The algorithm interpreted what the patient was trying to say with 75% accuracy.

Although the experiment was only conducted in one patient and only included asking the patient to try to say up to 50 words, the study shows that "the critical neural signals [for speech production] exist and that they can be leveraged for this application," said Vikash Gilja, an associate professor at the University of California, San Diego, who was not involved in the study.

To Moses and his team, this study represents a proof of concept. "We started with a small vocabulary to prove in principle that this is possible," he said, and it was. "Moving forward, if someone was trying to get brain surgery to get a device that could help them communicate, they would want to be able to express sentences made up of more than just those 50 words."

STAT spoke with Moses to learn more about the development of the technology and how it could be applied in the future. This interview has been edited for length and clarity.

What problems were you seeking to address?

It's kind of easy for us to take speech for granted. We have met people who are unable to speak because of paralysis, and it can be an extremely devastating condition for them to be in. It hadn't been understood before if the brain signals that normally control the vocal tract can be recorded by an implanted neural device and translated into attempted speech.

Can you describe how the technology works? What information goes

into it and how is that analyzed to produce words?

This is in no way mind reading; our system is able to generate words based on the person's attempts to speak. While he's trying to say the words that he's presented [with], we record his brain activity, use machine learning models to detect subtle patterns, and understand how those patterns are associated with words. Then we use those models with a natural language model to decode actual sentences when he is trying to speak.

What's the importance of including the natural language model?

You could imagine when you're typing on your phone and it figures, "Oh, this might not be what you want to say," that can be very helpful. Even with the results that we report, it's imperfect. It helps to be able to use the language model and the structure of English to improve your predictions.

What was surprising to you about what you learned during this study?

One of the very pleasant surprises was that you were able to see these functional patterns of brain activity that have remained intact for someone who hasn't spoken in over a decade. As long as someone [can imagine] producing the sounds of what their vocal tract would normally do, it's possible for us to be able to record that activity and identify these patterns.

How did you feel when you realized that the system was in fact producing the words that the patient was trying to say?

My first thought was "OK, that's just one sentence. It could have been a fluke." But then when we saw that it was working sentence after sentence. It was extremely thrilling and rewarding. I know that the participant also felt this way, because you can tell from looking at him that he was getting very excited.

What are the next steps to improving the system?

We need to validate this in more than one person. And we want to know how far this technology can go. Can this, for example, be used to help someone who's locked in completely — who only has eye movements and cannot move any other muscles. If we show that it can work reliably in people with that level of paralysis, then I think that that's a strong indicator that this is really a viable approach.

How do you envision this technology being applied in the future?

The ultimate goal really for us is to completely restore speech to someone who's lost it. That would mean any sound someone wants to make, the system is able to produce that sound for them by synthesizing their voice. You could even restore some personal aspects of the speech, such as intonation, pitch, and accent. It's going to be a lot of effort and we have a lot of work to do, but I think this is a really strong start.

STAT

Q&A: Microsoft research head explains how generative AI could help doctors be more human

By Nicholas St. Fleur | **MAR. 5, 2024**

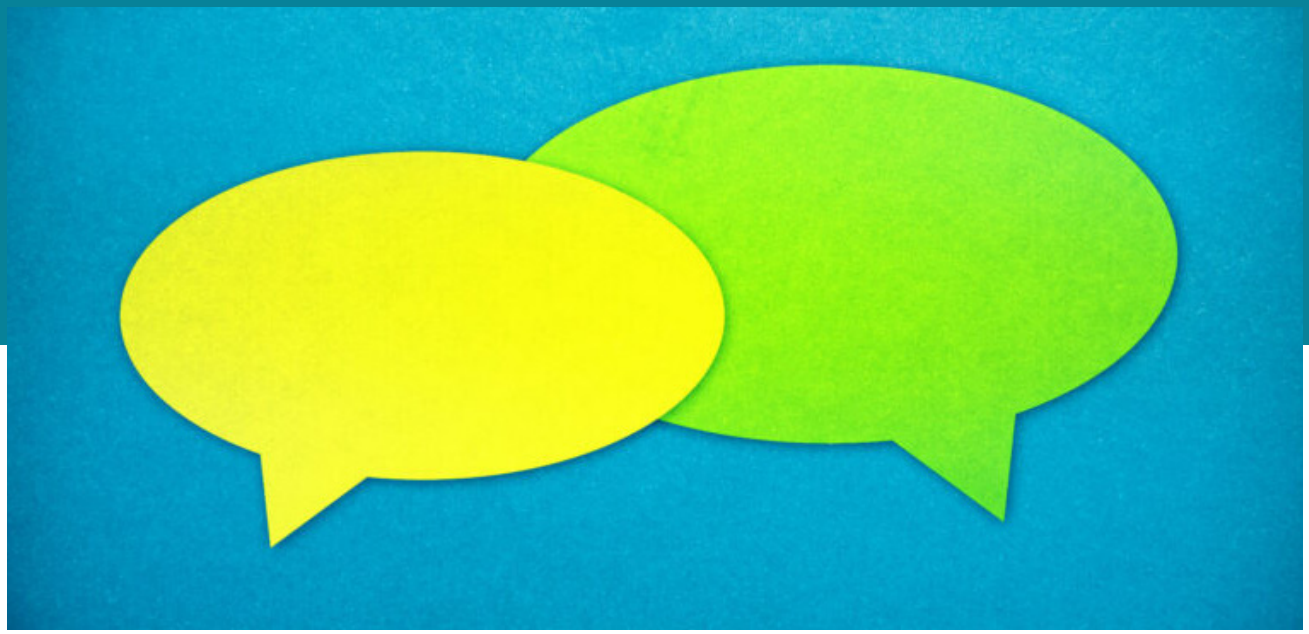

ALEX HOGAN/STAT

Generative AI is having a moment. Peter Lee, who oversees Microsoft's approach to health care and co-authored the 2023 book "The AI Revolution in Medicine: GPT-4 and Beyond," calls it "the most transformative tool ever developed in all aspects of health care and medicine."

Microsoft has been collaborating with the organization OpenAI for several years, so when OpenAI released ChatGPT in November 2022, Lee quickly realized there was a need to educate the medical world on the technology.

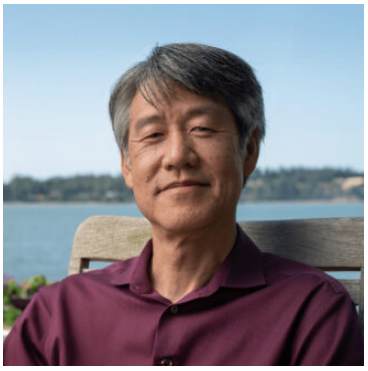

"I started to get emails from doctor friends of mine around the world, more or less all saying the same thing: 'Wow, Peter, this is amazing stuff and I'm using it in my clinic for this purpose,'" Lee told STAT. "That was frightening. … That started us down the long road of deeply investigating the benefits, as well as the limitations, of these models in the medical domain, as well as writing educational material, including a whole book written explicitly for doctors and nurses."

Lee — one of the 50 influential people named to the 2024 STATUS List — sat down with STAT to discuss his thoughts on how the technology can help physicians with tasks like after-visit summaries, as well as concerns about how it may perpetuate bias.

So much of my headspace, and the headspace of a lot of computer scientists, is on generative AI. That's due to our collaboration with OpenAI. You have this big AI model that's been trained to do exactly one thing, which is predict the next word in a conversation. That's all that it has been optimized for. It [has] literally not been optimized for anything else.

Like ChatGPT in this case?

Yep! And yet, to do that really, really well, the machine learning system has had to discover how to do arithmetic. If you want to pick the perfect next word in the conversation, if someone says, "2+2 = 'blank,'" the best way to answer that question is to actually discover how to do addition. Or to predict the next word in a conversation where the last sentence is "… and the killer is 'blank.'" To be able to solve a problem like that, you have to be able to read an entire murder mystery and do all the deductive logic to do that simple thing in next word prediction.

What's so beautiful is that we think of next word prediction as just being this trivial thing, sort of like the autocomplete on your iPhone. But when you do it at this gigantic, unimaginable scale, the machine learning system has had to self-discover how to do all this thinking, and there's just something that I find so incredibly beautiful about that. And that even then maps to medicine and health care.

How so?

Let's have a conversation: You're my patient. I'm a doctor. We're talking back and forth. I signed you up for labs, and I get a lab report and then the last sentence in the conversation is, "… and the diagnosis is 'blank.'" To optimally predict that next word in the conversation implies that the machine learning has to learn medicine. Isn't that wild? I think it's the most surprising, astounding, and beautiful thing in the world today.

What are your bug concerns with generative AI being used in the medical field?

The main thing I've been trying to teach doctors and nurses is that

if your mental model of a computer is that it's a machine that does perfect memory recall and perfect calculation — then the most important thing to understand about generative AI is that it's not a computer. It's a reasoning engine or a thinking machine, but it also has some of the same limitations as a human brain.

If you ask it to regurgitate something by rote, it might hallucinate because it can't remember it perfectly. If you ask it to do a big, complicated math problem, it might get it wrong in the same way that a human can. At the same time, it can opine in incredibly sophisticated ways about connections between concepts.

Sticking with this theme of finding beauty in what seems complex, what is so beautiful about GPT-4 and its applications in medicine?

Maybe two things. The most surprising one has been GPT-4's ability to grasp what psychologists call "theory of mind."

One of Epic's uses of GPT-4 is to help doctors write "after-visit summaries" to their patients. So you're my patient, you come to see me, and then you get sent home. Then I have to send you an email with instructions on how to take care of yourself after I've treated you. After-visit summaries are a pain in the ass for doctors to write because they have to get them right. They have to access four or five parts of the electronic health record to make sure that they're getting the prescriptions and the other things to do at home. The chance of a malpractice suit if they get this wrong is extremely high.

Epic has integrated GPT-4 to go grab all that information and then draft a note for the doctor to review before sending it out. In the

early tests, these are controlled clinical studies, patients are rating the GPT-4-written emails as more human than the ones written by doctors. It's not the case that an AI is more human than a doctor. It's just the opposite, obviously. But the AI has the tireless ability to pick out [personal information from] all of the health records and all the transcripts of the conversations [and then] put in that extra little line like, "Congratulations on just becoming a grandparent," or, "Best wishes on your daughter's wedding in Maine next month." Those extra little touches where AI adds that personal touch actually makes a meaningful difference in the patient experience.

The second thing is the general intelligence of GPT-4. We find when it's trying to make a medical diagnosis, it's able to triangulate between multiple specialties. So if you come to me feeling anemic, it could be a problem for an endocrinologist, or a cardiologist, or a nephrologist, or for a psychologist. Depending on what doctor you go to or what specialist your primary care physician refers you to, you're going to get different diagnoses. GPT-4 is able to look at your condition, your labs, and your initial presentation from all those perspectives at once. What we see consistently is that by doing that, it's coming to more well-rounded assessments.

That's fascinating. So can generative AI help human doctors be more human?

It's such a good question. We talk about humans prompting the AI, but there are times when the AI can prompt the human to just take a step back and reflect, just for a moment longer, about a potentially difficult situation. Put yourself in the shoes of the doctor. [Epic using GPT-4] proposes a note that puts at the end, "Congratulations on your son's

high school basketball team winning the championship!" It actually causes the doctor reading that draft, maybe for just three seconds longer, to reflect on the life of that patient. It's what I've been calling a "reverse prompt."

It sounds beautiful and creepy and wonderful and disturbing all at the same time.

As we're marching towards this future where generative AI is in doctors' offices and operating rooms, how do we deal with concerns over privacy and bias?

Privacy, I don't think is an issue. The OpenAI services on Microsoft's Azure cloud provide HIPAA compliance. We're very proud of that but our cloud is not unique. AWS and Google Cloud provide the same sort of compliance to enterprise customers of those clouds. That's different from the consumer space. If you're using a consumer product like Google Search or ChatGPT the privacy guarantees aren't as strict. But if you're a healthcare organization that subscribes to Microsoft Azure, you get HIPAA compliance.

Bias, that's a serious issue and a potentially devastating one because these models are, in my view, hopelessly biased. They're hopelessly biased because they learn from us and they are hopelessly biased in the same way that human beings are hopelessly biased. There are computer scientists, colleagues of mine, who believe that we can fix these things, but I don't believe it. The thing to understand is that while they are hopelessly biased, these AI systems also understand the concept of bias and why it's bad.

One of the things we find in our research is that if you describe a situation that involves a decision about a patient and you give it to GPT-4 and ask it to check for biases, it can outperform human beings in identifying biased decision-making. In an experiment with the New England Journal of Medicine, we were having GPT-4 read submitted manuscripts. Consistently, GPT-4 was able to spot non-inclusive language and bias in the manuscripts that escaped the attention of the human reviewers. There's a potential here that generative AI can be one of the most powerful tools in combating bias, even though those tools themselves will be prone to making biased decisions in the same way that humans [are prone to].

I would not trust an AI system on its own to make decisions on whether my health insurance claim should get reimbursed or not. I think a human being should be on the line to do that. But I would like a generative AI system to be a second set of eyes to check whether that human being that's deciding my insurance is being biased against me or not, because I think the AI system can be really good at that.

STAT

Nvidia says generative AI will revolutionize health care. So did IBM, with Dr. Watson

By Casey Ross | **MAR. 25, 2024**

MIKE REDDY FOR STAT

The health care sector is a notorious laggard when it comes to technology. It was slow to use computers, digitze patient data, and share information electronically. While most of the world instantly interacts viz Zoom and Slack, hospitals — even today — are still sending faxes.

But something different is happening with generative AI.

Health systems, drugmakers, and insurers are racing to build the technology into their operations, aligning themselves with corporate giants such as Microsoft, Google, and Nvidia, whose executives speak about each incremental advancement of AI as an earth-shattering event. "The generative AI revolution is here," Jensen Huang, 61, Nvidia's leather-jacket-clad CEO, declared during the company's GTC conference in San Jose, Calif., last week.

That may be so. The technology is undeniably getting more powerful, and more beneficial to businesses in health care and beyond. But many physicians, chemists, and computer scientists also fear that the overheated rhetoric — and the thinly veiled profit motives behind it — will inevitably result in yet another letdown.

"One way that people lose faith in science and medicine and institutions is that they're promised things that are not delivered," said Rob Patro, a computer science professor at the University of Maryland. The marketing of AI in health care is particularly pernicious, he said, because it convinces investors, patients, and practitioners that changes likely to take decades are achievable overnight. All one needs to do is believe, and make a bigger bet now.

"There is a huge difference between making a scientific discovery and making a change in health care for everyday patients," Patro said. "I don't think it's productive scientifically to think of everything as transformative, or a breakthrough, or revolutionary."

And yet these are the words that drive the marketing of AI. Each new

microchip is a "transformation." Every AI model — to transcribe clinical notes, answer patient emails, or predict a protein's shape — is a "revolution," as if the distance between a new piece of technology and a cure for a terminal illness is no longer than a 500-word press release.

Health care learned this same lesson a decade ago, when IBM overhyped its Watson supercomputer during the deep learning run-up as a cancer cure, only to have its health care business collapse under the weight of its own exaggerations.

Much like IBM, the purveyors of generative AI are now promising to transform everything at once: the discovery of drugs, the treatment of patients, and the collection and analysis of massive amounts of data. Microsoft asserted in a recent blog post that, based on a company-commissioned survey, health care businesses that adopt AI get an average return of $3.20 for every $1 they invest. The post cited additional survey data to show evidence that its AI tool for automating documentation of patient visits is easy to use and saves time.

But in health care, company-sponsored survey data is not enough evidence to establish that a technology is effective, that it works equally well on different patients, or that it will save money and improve care in the long run.

Getting to those results requires digging deeper into AI's use within individual hospitals — a process that takes years and involves dozens of discrete tasks to evaluate a product's validity, safety, and reliability, according to the Health AI Partnership, a research group led by Duke University that advises hospitals on AI adoption.

A parallel process must play out within drug discovery, in which generative AI's ability to model disease-causing proteins and identify promising drug candidates still must be verified through clinical trials, which fail at a rate of 80% due to lack of efficacy or toxicity. "Fixing that with AI? Yeah, that's really, really going to be hard, because of the sheer lack of knowledge in these areas," said Derek Lowe, a medicinal chemist and author of "In the Pipeline," a blog on drug discovery.

But those kinds of details tend to get lost on corporate stages where AI is a magic act.

Aside from OpenAI, the maker of ChatGPT, no company has benefited more from the technology's rise than Nvidia, which makes the microchips and supercomputers that power AI models. Its stock price has jumped more than 230% in the last year, making it one of the most valuable companies in the world, with investors pegging its worth at more than $2 trillion.

Given the ubiquity of its microchips, the company's fate will not be determined by its performance in health care alone. But health care is a domain in which Nvidia has been making significant inroads, signing partnerships with hospitals and major drug and device makers such as Amgen and Johnson & Johnson.

"All the digitization we've done for genes and proteins and amino acids — that digitization capability is now passed through machine learning so that we understand the language of life," Huang said during his talk in San Jose. He unveiled a product that provides easy access to generative AI models to help discover new drugs and interact directly with patients to book appointments and carry out other tasks.

But it remains unclear whether these AI models will anytime soon result in new medicines or better care for patients.

"What's going on with Nvidia — it's hard to know if it's real or not," said Howard Forman, a radiologist and professor of health management at Yale University. He said the rise in its stock price, while an impressive measure of enthusiasm, has also created unrealistic expectations. "Most people would say it's a bubble financially because it's almost impossible to imagine a growth rate that could meet this," Forman said.

Nvidia is not alone in aggressively pitching AI. Epic Systems, the nation's largest electronic health record vendor, is racing to develop dozens of generative models. Microsoft, Oracle Cerner, Amazon, and Google are also making huge investments. Insurers and academic health systems, meanwhile, are eager to be part of their revolution if it means offloading burdensome administrative tasks. Many are incorporating AI tools to automate clinician note-taking, answer patients' questions, and draft summaries of hospitalizations, radiology reports, and other documents.

Executives involved in these deals say they've never seen health care move so quickly to embrace a new technology. "The builders of all these amazing tools and the users are coming closer and closer together in terms of how they're deployed," said Richard Clarke, chief data and analytics officer at Highmark Health, a Pittsburgh-based insurer and health services company that uses AI tools provided by Google's cloud business. "It brings the research to practice faster."

But keeping that momentum going will mean overcoming history and the hurdles that tend to make health care problems so intractable.

Developing new medicines and changing entrenched work routines is not as simple as downloading new software.

Forman thinks back to 1988, when he was a medical student applying for residency. Every program told him about this dramatic change unfolding: a new electronic record-keeping system was about to transform data collection and the delivery of care.

It was only 25 years later that Yale, the place where he now works, completed that transition.

"That was sobering for me," Forman said, adding that, despite the experience, he remains optimistic about AI's ability to drive change. "I'm just not sure we're at an inflection point, per se, even though the rest of the world thinks we should be."

This story is part of a series examining the use of artificial intelligence in health care and practices for exchanging and analyzing patient data. It is supported with funding from the Gordon and Betty Moore Foundation.

SPONSORED CONTENT

IMO Health

NLP and generative AI in life sciences and precision medicine

By IMO Health

Using linguistics, natural language processing (NLP) tools, deep learning, and technology to extract meaning from text is powerful yet challenging. It requires understanding words themselves (lexicography), how words form a sentence (syntax), the meaning of words and texts (semantics), synonyms, abbreviations, and more. Using machines to understand meaning in clinical and biomedical text is far more complicated, requiring semantic tools to identify complex concepts and relationships between diseases, drugs, genes and variants, cells, pathways, compounds, proteins, biomarkers, and phenotypes. most current NLP methods must be trained on relevant data, and even NLP based on generative AI such as ChatGPT, must be fine-tuned to reliably answer questions such as:

- What is the biomedical mechanism for a particular disease?
- What are similar diseases to a known disease which may be treated by a certain drug and share common signaling pathways?

- Are there "treatment" relationships between one known drug with diseases that are outside of original indications (off-label usage)?

In this paper, we review applications of NLP and generative AI in life sciences and precision medicine, exploring the nuances and complexity that demand specialized NLP models and biomedical domain expertise when extracting insights from medical literature, trial protocols, patient data, real-world evidence, drug labels, and post-market surveillance.

Drug Discovery and Drug Repurposing

A critical step in drug discovery and drug repurposing is identifying evidence from massive and rapidly growing biomedical literature to help generate hypotheses. The use of systematic literature review and data mining support this work to build a knowledge base, assist research gap analysis, synthesize evidence, and direct research. Effective literature review also includes details that support traceability requirements for FDA regulatory submission.

The process of literature review poses many challenges. It is:

- Labor intensive due to the volume of articles in sources such as PubMed Central (PMC), MEDLINE, and Online Mendelian Inheritance in Man (OMIM)
- Prone to errors as a highly manual process
- Difficult to stay current with sources that grow and evolve rapidly

Generative AI and NLP for Drug Discovery and Drug Repurposing

Generative AI can significantly improve literature analysis for drug discovery and drug repurposing. The combined use of NLP and generative AI supports each step of the process from study protocol setting and literature retrieval, to abstract screening, full-text screening, data element extraction from full-text articles, results summary, and data visualization. Unique NLP tasks predict articles' relevance based on their title, abstract, and other metadata. Named entity recognition parses full-length articles and extracts data elements from both etxt and tables and highlights supporting information. With AI-automated literature review and mining, one can specify protocol with natural language, speed the process to define fine-tuning tasks, and create a "living" system that proactively and continously updates relevant literature in a timely manner.

While the steps in each literature review are the same and can be more efficient with AI, the NLP requirements to address different business goals, diseases, and compounds vary in different contexts. Even work in the same disease space requires unique knowledge to produce reliable results. When augmented with domain expertise, data scientists can fine-tune and specify a protocol to customize and refine models that extract and summarize information unqiue to each study. Ultimately, scientists can dedicate more time to ensuring data quality and synthesizing evidence while staying current.

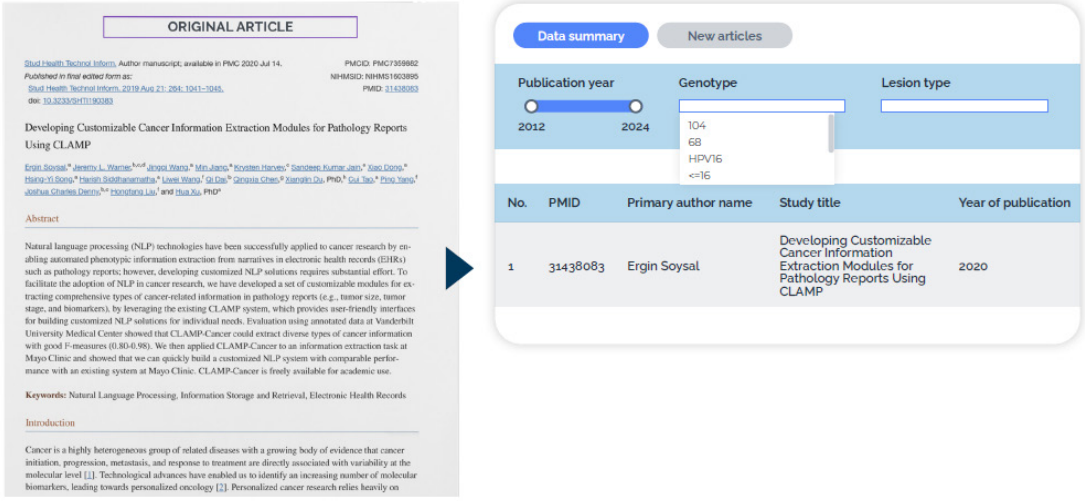

Unique NLP tasks predict articles' relevance based on their title, abstract, and other metadata.

Soysal E, Warner J, Wang J, et al. Developing Customizable Cancer Information Extraction Modules for Pathology Reports Using CLAMP. Stud Health Technol Inform. Aug 2019. Accessed via: https://pubmed.ncbi.nlm.nih.gov/31438083/

An IMO Health Example:

Accelerating Drug Repurposing with an AI-Driven Framework

A recent study developed a framework to apply generative AI for drug repurposing studies. IMO (Melax Tech) scientists used NLP to extract biomedical entities and relations from 35 million PubMed abstracts. Using deep learning-based models, they built a knowledge graph of 20,000 entities (drugs, diseases, genes, etc.) and 10 million relations ("inhibits," "treats," "stimulates," etc.) and scoring systems to predict the "treats" relations for each drug-disease pair. The evaluation module applied link prediction for 15 successful pairs of drugs and their new indications and found that all are ranked in the top 0.5% across all diseases.

Clinical Trial Optimization

Clinical trials are the gold standard used to evaluate the effects of vaccines, drugs, medical devices, and treatments on human health outcomes, assessing the benefits and harms vs. standard treatments. Unstructured, narrative text is at the heart of several key trial steps, including protocol definition, parsing clinical trial protocols, and patient recruitment.

Recruiting a representative and clinically meaningful population is a crucial step and one of the biggest barriers for the successful implementation of clinical trials. Suboptimal criteria selection can lead to low accrual, resulting in trial incompletion. Overly rigid criteria restrict patient access and may reduce the potential relevance for patients that could otherwise benefit from the intervention. Challenges to effective and efficient recruiting include extracting eligibility criteria from lengthy clinical trial protocol documents and creating clear, succinct criteria. This task is critical, yet time-consuming and inefficient.

Generative AI and NLP for Clinical Trials

Integrating NLP and generative AI into clinical trial design and recruitment reduces the time required to initiate and conduct clinical trials, enhances the representativeness of the participant pool, and supports stronger trial results, ultimately benefiting both investigators and patients.

To standardize eligibility definitions, NLP techniques can automatically extract criteria from clinical trial protocol documents,

identifying disease cohort characteristics, summarizing them into variables, and extracting endpoint measures. When NLP is optimized in medical domains and across disease areas, domain experts can fine-tune the approach using prompts to specify entities, demographics, lab texts, and biomarkers that are unique to the trial. They can get to specific relevant attributes with values, modifiers, and related conditions for both inclusion and exclusion.

Further, by structuring clinical trial information into a knowledge base, trial designers can simulate alternatives and optimize criteria. With links to real-world data, trial designers can more readily determine how many patients in a database are eligible.

NLP also creates efficiencies in patient recruitment. By generating criteria text and executing queries against medical records data in the EHR, it creates an electronic eligibility process to prescreen and validate patients.

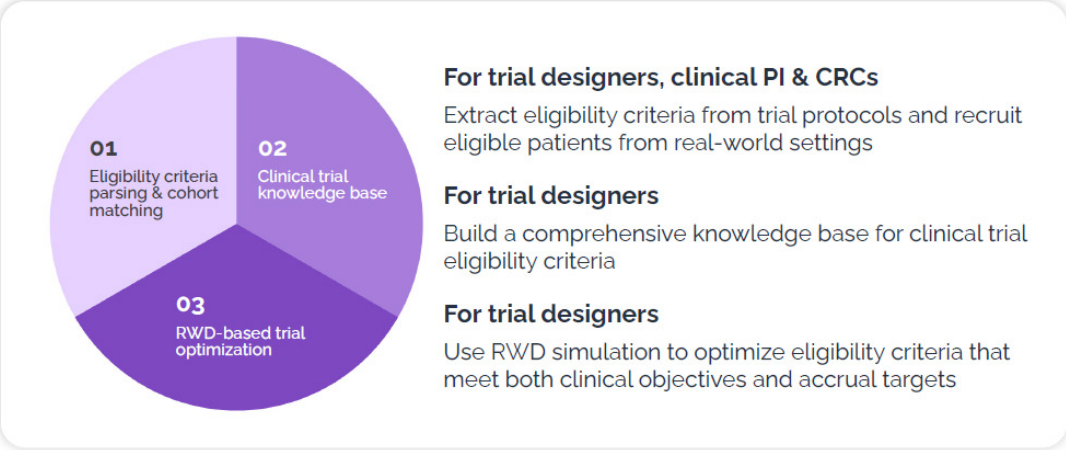

NLP and generative AI support trial designers across multiple steps of the process.

An IMO Health Example:

Advancing Clinical Trial Study with AI-Powered Eligibility Criteria Extraction

Using clinical trials data acquired from ClinicalTrials.gov, spanning oncologic, neurodegenerative, autoimmune, endocrine, and circulatory system disorders, IMO (Melax Tech) scientists developed a system comprised of pre-processing, knowledge ingestion, GPT-based prompt modeling, post-processing, and interim evaluation modules. the system evaluated 180 manually annotated trials covering nine distinct diseases and exhibited outstanding performance in criteria entity identification, consistent proficiency, and effective handling of the intricate contextual aspects of criteria. The results produced accuracy of 78.95% across a wide range of diseases.

An IMO Health Example:

Extracting Postmarketing Adverse Events from Safety Reports in VAERS Using Deep Learning

This study implemented and evaluated deep learning algorithms for named entity recognition to extract nervous system disorder-related adverse events. IMO (Melax Tech) researchers collected Guillain-Barré syndrome (GBS) related influenza vaccine safety reports from VAERS from 1990 to 2016 and manually annotated 91 reports with 2,512 entities related to nervous system disorders. The team evaluated a variety of conventional machine learning and deep learning algorithms against a domain-specific Bidirectional Encoder Representations from Transformers (BERT) model pre-trained using VAERS

reports. Deep learning-based methods outperformed conventional machine learning-based methods. An ensemble of the BioBERT and VAERS BERT models achieved the highest exact match micro-averaged F-1 score at 0.6802 among peer models.

Disease Phenotyping and Precision Medicine

Understanding the patient journey for a particular disease is a powerful way to help assess the benefits, harms, and trajectory of medical treatments, deliver precision medicine, and improve outcomes. Massive amounts of detailed patient data are available in EHRs to support this analysis, but must be accurate, structured, and well-organized to unlock its value. Further, conducting meaningful research on patient populations requires assembling patient cohorts with similar disease and treatment profiles.

Challenges to creating and assessing patient journeys start with inconsistencies in EHR data and real- world evidence. With inconsistent implementation of interoperability standards, data becomes highly variable and lacks specificity, creating gaps that require a great deal of manual intervention and can drain clinical resources. In addition, free text data in clinical notes including pathology, radiology, and radiation therapy reports often require deep domain expertise to extract meaning.

Generative AI and NLP for Disease Phentoyping and Precision Medicine

NLP and generative AI are ideal tools to extract clinical information from EHRs, however, general healthcare NLP models often fall short.

To be effective and keep up with a knowledge base that grows and changes rapidly, NLP solutions must be trained on data that is optimized in medical domains. Models must understand and extract all possible genes, diseases, variants, and mutation patterns, identify disease associations between phenotype and variant, as well as characterize rare diseases and variants. Solutions must be able to normalize concepts to standard ontologies such as MedDRA and Medical Subject Headings (MeSH) and use pattern recognition to identify complicated categories of information such as diseases and symptoms caused by various gene-protein mutations.

With specialized clinical and biomedical NLP and generative AI, domain experts can use prompt engineering to extract insights from clinical notes and specialized reports, finding patient characteristics on disease progression, trajectory, treatment, and procedures, as well as extract dates associated with each. Understanding the journey across a population of patients is more complicated. It requires compiling and aligning data across multiple patients to create real-world evidence journeys. To compare disease and treatment progression, every patient needs a clear, age-based timeline from birth through each key milestone – the disease, symptoms, and conditions. With aligned timelines, researchers can evaluate progression and outcomes, and trigger screening and biomarker testing for individual patients.

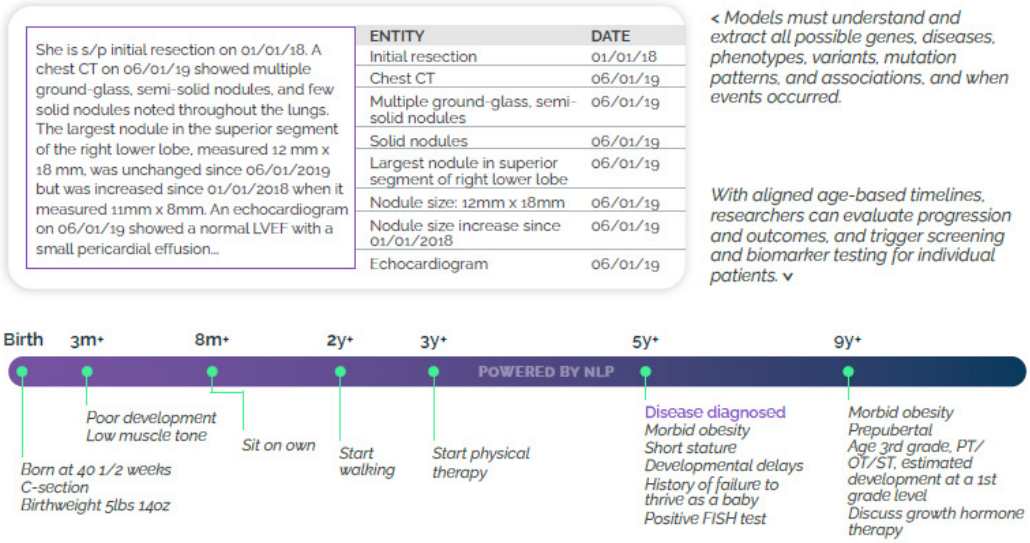

An IMO Health Example:

Extracting Crucial Treatment Details from Radiation Oncology-Specific EHRs

In a recent study, IMO (Melax Tech) scientists used NLP to extract free-text data from radiation oncology-specific EHRs, developing customizable modules for cancer-related information in pathology reports including tumor size, tumor stage, and biomarkers. Based on data elements suggested by the College of American Pathologists, the study used 400 randomly selected pathology reports from cancer patients. For named entity recognition, it implemented regular expression-based, dictionary lookup- based, as well as machine learning-based approaches. For relation extraction, it developed rule-based, machine learning, and hybrid approaches. When evaluated against existing systems, the customized NLP pipeline achieved comparable performance with reduced production time and greater adaptability.

Conclusion

Advancements in life sciences and precision medicine must tap into abundant sources of free text documents to inform decisions. When supported by domain-optimized NLP models and biomedical experts, NLP and generative AI solutions are creating efficiencies that account for nuances and complexity to unlock reliable insights that can accelerate and enhance drug discovery, clinical trials, and post-marketing surveillance, and improve patient care.

To learn more about how IMO solutions can unlock deeper insights in life sciences applications, visit imohealth.com/life-science-and-clinical-research.

References:

- Huang LC, Li Y, Lee K, et al. Knowledgesphere: An Automated and Integrative Framework for Drug Repurposing Empowered By Knowledge Graph and AI. Value in health, Volume 26, Issue 6, S2. June 2023. Accessed via: Link
- Data S, Lee K, Paek H, et al. AutoCriteria: Advancing Clinical Trial Study With AI-Powered Eligibility Criteria Extractions. Value in health, Volume 26, Issue 11, S2. December 2023. Accessed via: Link
- Du J, Xiang Y, Sankaranarayanapillai M, et al. Extracting Postmarkerting Adverse Events from Safety Reports in the Vaccine Adverse Event Reporting System (VAERS) Using Deep learning. Journal of the American Medical Informatics Association, Volume 28, Issue 7. July 2021. Accessed via: Link

- Soysal E, Warner J, Wang J, et al. Developing Customizable Cancer Information Extraction Modules for Pathology Reports Using CLAMP. Studies in health Technology and Informatics. Volume 264. August 2019. Accessed via: Link

STAT

AI and rural health care: A paradigm shift in America's heartland

By Bill Gassen | **JUNE 26, 2024**

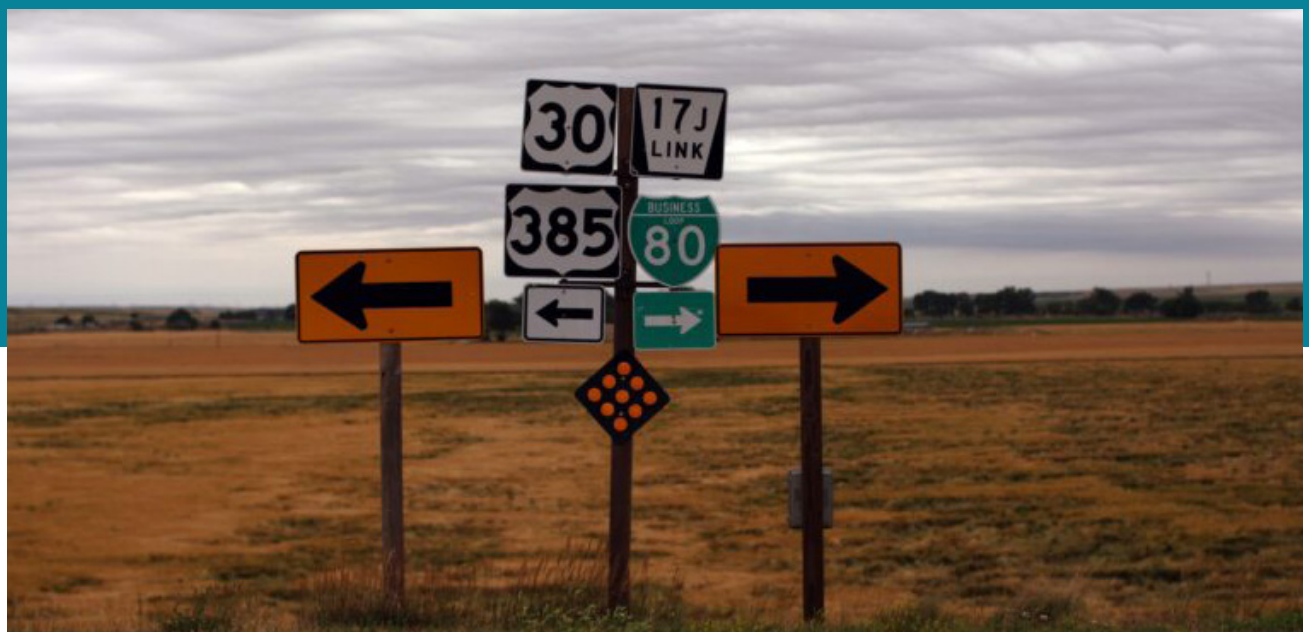

SPENCER PLATT/GETTY IMAGES

The use of artificial intelligence is happening in "unlikely places." So says a recent working paper published by the National Bureau of Economic Research. While much of the nation is debating the future of AI, health care providers in rural America are pioneering new uses of it in their practices. As the leader of the U.S.'s largest rural health

care system, I predict the impact of AI on rural health care will be consequential.

After a decade of decline before the pandemic, a recent study by the U.S. Department of Agriculture indicates that the population in rural areas is rising a bit. Rural counties across the country — defined as those with cities of up to 50,000 people — grew one-quarter of a percent from 2020 to 2022.

That tiny population surge, however, isn't likely to mend the greatest demographic challenge for rural health care: recruiting and retaining enough clinicians to work in the medical specialties that are in short supply across rural America.

Three out of five federally designated medical provider-shortage areas are in rural regions. Rural communities have only 30 specialists per 100,000 people, compared to 263 specialists per 100,000 in urban regions. Meanwhile, 25% fewer rural physicians will be practicing by 2030 due to an aging workforce with looming retirements.

A paradigm shift

Nearly all of America's doctors are experiencing burnout as the pace of their practice has relentlessly sped up in recent years. But many are encouraged by the potential for AI to help them improve efficiencies in ways that allow them to refocus on their patients, rather than trying to keep up with electronic health records.

A paradigm shift is happening in rural America as rural health providers come to embrace the idea that what we do won't change, but *how we do* it must.

Clinicians working in Nebraska corn country at Bryan Health are now using AI-enabled software that takes notes on their phones with the press of a button so they can look their patients in the eyes instead of clicking away at a keyboard. The technology securely records a provider's conversation with a patient during a visit using "ambient listening," which is then transcribed in the electronic medical record.

Bryan Health CEO Russ Gronewold told me that one physician in Grand Island called this "career extension technology." Another in Lincoln said he'll look back at this moment as "one of the most pivotal moments" in his medical profession.

In the next few weeks, Bryan Health will also launch a new generative AI tool in its electronic medical record that's designed to reduce the significant amount of time physicians spend responding to large volumes of patient messages. The new tool will "pre-populate responses," but physicians will have the ability to tailor and edit each message before sending.

While recent studies have shown that the use of large language models may not actually save clinicians time, saving time might not be the only measurement that matters. As Dr. Michael Pfeffer, the chief information officer and associate dean at Stanford Health Care and Stanford School of Medicine recently shared, his team found that generated draft responses for patient messages "reduced cognitive burden" by giving physicians a place to start.

In the dairy country of Wisconsin and the upper peninsula of Michigan, Marshfield Clinic Health System is taking a slightly different approach when it comes to combating physician burnout related to

patient messages. Marshfield will soon deploy AI technology within its electronic medical record to "reduce noise" for physicians by sorting and routing messages to the appropriate member of the care team.

According to Marshfield's chief information and digital officer, Jeri Koester, nearly 60% of the system's patient-initiated messages are related to prescription refills, scheduling, or completing forms — tasks that can be managed by a nurse or medical assistant. This new tool will "remove clutter" from physicians' inboxes and allow them to instead focus on clinical and urgent messages that require their expertise and immediate attention.

When the U.S. Preventive Services Task Force recommended colorectal cancer screenings begin at age 45, clinical teams at Sanford Health gathered to determine how to manage screening for 100,000 newly eligible individuals in the rural Dakotas, where there was a limited supply of gastroenterologists. They developed their own AI model that includes additional risk factors that may put people at heightened risk for colon cancer. The model is designed to help physicians understand the risk of the patient in front of them without having to scroll through medical records, saving doctors time they can spend with the patient instead.

AI's next chapter

AI can and will do much more than streamline administrative tasks. These technologies will soon serve as another tool in clinicians' black bags — wherever their practices are located — metropolis, suburb, or farming town. AI-enabled clinical decision-support tools will help to identify serious health threats, improve diagnoses and calibrate the

precision of medical treatments.

A new effort, led by the White House, is focused on developing a voluntary framework of health care AI commitments. The initiative does not shy away from the trust barriers that must be addressed, both for consumers and health care practitioners, including ensuring data that an AI model is trained on is representative of the population it will serve as a guard against bias.

So far, 38 payers and health systems have come together in this collaboration to determine how to harness AI models safely, securely and transparently. Bringing diverse voices to the table is a critical component of this work.

The challenges are not the same for urban and rural health care systems.

Rural America has some of the highest rates of late-stage breast cancer diagnoses in the country. A recent report from the Centers for Disease Control and Prevention found that rural Americans are more likely to die early from preventable causes like cancer. The potential to reverse this trend through new AI technologies that forecast the risk of disease will be a game-changer for rural clinicians and the patients they care for.

Preterm birth is another rural issue. In rural northern Minnesota, one of the poorest and most geographically isolated regions in the state, OB-GYNs at Sanford Bemidji Medical Center launched a pilot using FDA-cleared AI-enabled non-stress test belts to monitor fetal heart rate and the presence of contractions among patients who may be at higher

risk of pre-term delivery, allowing them to intervene earlier to ensure the best possible maternal health outcomes.

My children are enamored with a movie called "The Croods," which tells the story of a family of prehistoric cave-dwellers about as far from high-tech living as one could imagine. In the original 2013 film, the stubborn, cautious father ("Fear keeps us alive!") refuses to let anyone leave the cave except for brief forays at daybreak to gather food, admonishing his children to "never not be afraid." But they defy that edict and eventually make it to the other side of the mountain to see a peaceful paradise there.

It makes sense to be cautious about AI in health care, no matter where one lives or practices medicine. But some health care providers, including those in the most remote and rural locations in our nation, have already crept over the mountain and have seen a new world of promise.

Bill Gassen is president and CEO of Sanford Health, the largest rural health system in the United States, with headquarters in Sioux Falls, South Dakota.

STAT

Could billions of dollars in AI funding lead to the same number of new drugs?

By Ashu Singhal and Sajith Wickramasekara | **JUNE 25, 2024**

ADOBE

Drug discovery has quickly become the most enticing place to apply artificial intelligence. Billions of dollars are being invested in AI-driven "techbios." In an industry where nothing changes overnight, even large biopharma companies are touting AI as key to how they're transforming their discovery engines.

But in the race to integrate AI into drug discovery, investing so heavily in scaling one part of the system overlooks the rest. Failing to reimagine R&D systems to handle the new speed and scale of AI-driven discovery risks overpromising and under-delivering to the people who need new medicines.

There's no question that new AI models for drug discovery deserve serious attention. Within the next five to 10 years, AI will fundamentally change the way drugs are designed, with the potential to produce an order of magnitude more high-quality candidates against a broad range of new diseases. In the last year alone, AI has been used for identifying novel targets in areas like cardiomyopathy, generating novel antibodies, and even designing newer modalities like optimized mRNA vaccines for influenza.

But the focus for AI can't just be on discovery. The rest of the pharma R&D system — how drugs are developed, tested, approved, and manufactured — will need to accommodate this vast increase in speed and scale.

Pharma R&D has been heading in the wrong direction, becoming progressively less efficient over the last few decades. Large pharma companies now spend more than $6 billion on R&D per approved drug, compared to just $40 million (in today's dollars) in the 1950s, with approximately 85% of that spending coming after discovery. The inequality between near-zero-cost AI-driven drug discovery and skyrocketing costs for clinical trials and regulatory approval will create a bottleneck, stalling promising drugs from reaching patients unless AI is equally applied across the entire R&D lifecycle. As co-founders of Benchling, a technology company focused on life science R&D, we've

heard about this issue from scientists, R&D leaders, and CEOs across more than 1,200 companies.

To be sure, the biopharma industry has long known it needs to evolve its R&D systems. The paper describing the rapid decline in R&D efficiency, often known as Eroom's Law, is more than a decade old. AI, now being used across nearly every industry as a force for disruption, bringing with it automation, scalability, and intelligence, should be used to improve every part of the R&D lifecycle, not just discovery, to increase throughput and improve efficiency.

Rethinking the R&D lifecycle with AI

This is no small lift — applying AI will involve rethinking how R&D organizations operate, rather than simply applying AI to how things work today.

Just look at how AI-driven drug discovery is already putting new and different pressures on lab-based experimentation. Drug candidates discovered with AI need to be tested through experiments in the lab, which in turn generate experimental data that is used to further refine AI models through a "lab in a loop" process. An enormous influx of new AI-generated candidates and new data-hungry AI models means labs need to run experiments at a vastly greater scale.

That can't be achieved by simply optimizing the manual processes at the bench that labs rely on today. Instead, labs need to be reinvented around complex imaging and single-cell omics assays that allow a more complete understanding of biology, along with robotic automation that enables these assays to be run at scale. Although this

trend has already started, AI applications can rapidly accelerate it.

The dearth of specialized engineers needed to analyze data from complex assays and build robotic orchestration is a key bottleneck. AI models like scGPT can accelerate data analysis by automating code-intensive tasks such as reference mapping or cell annotation. AI agents will also enable scientists to set up robotic automation through natural language, democratizing access across the industry.

Advancements in AI also have the potential to address key challenges in clinical trials. Take, as an example, patient recruitment, the most time-consuming part of a trial. Even though millions of people are needed to participate in clinical trials, fewer than 5% of Americans have participated in clinical research of any kind.

Sound clinical trial design requires randomization, but that takes participants out of the driver's seat — they no longer have final say over their treatment decisions, creating a barrier to recruiting. In 2022, the European Medicines Agency provided a qualification opinion allowing the use of AI models to develop predicted control outcomes for Phase 2/3 trials from historical control data, ultimately requiring fewer participants to make this difficult randomization choice.

Beyond the lab work and clinical trials required to get a product to market, there is also so much knowledge work, from R&D managers reporting decisions at key program milestones to medical writers drafting filings for health authorities, quality assurance staff confirming data integrity, and much more. This knowledge work is about translating R&D data into decisions and documentation, and requires answering questions and generating content in the natural language of

scientists. Scientific large language models that are fine-tuned versions of popular general-purpose large-language models like GPT, such as BioGPT, or Llama, such as BioMedGPT-LM, have obvious potential.

But a generative pre-trained transformer built for scientific language isn't enough. The real challenge is rethinking how the underlying R&D data are structured and managed. To automate knowledge work, these large-language models need to operate on top of data that comes from the lab, a foundation in which there are many problems today. Large pharma companies often employ hundreds of software applications within R&D labs alone, leading to data silos and a lack of data standardization and interoperability that make it excruciatingly difficult to apply AI effectively. At Benchling, we're taking the same modern platform approach that has transformed how many businesses digitally manage their sales or financial data and applying it to R&D to make automation of knowledge work a reality.

AI will change how pharma competes

Another key element in the work of biopharma companies needs to be reinvented to make the promise of AI in biotech a reality: how companies compete.

Profluent Bio, a Berkeley, Calif.-based biotech, recently open sourced a novel, AI-designed, CRISPR-based, human gene editor. To open source such intellectual property was previously unthinkable. In an era where scientists working with AI can design many more drugs than we could possibly ever bring to market — imagine a world of drug abundance! — the competitive focus will shift away from protecting intellectual property and toward speed to market and creating step changes in the

efficiency of clinical trials and regulatory approvals.

This shift will, in turn, help solve the biggest issue in making AI-driven drug discovery even more powerful: access to data. The historic focus on intellectual property has created an industry culture that treats all experimental data as proprietary. Yet the success of AlphaFold2 and AlphaFold3 is entirely predicated on the public availability of protein sequences and experimentally-resolved structures. Progress in developing new foundation models will require data abundance.

Open-source software has long been a tenet of the tech industry, fundamentally altering the nature of collaboration and competition. If the biopharma industry wants to realize the benefits of AI, companies must work together to generate the data needed to power it. Companies are already starting to collaborate on open-source software projects around managing data in areas like molecular modeling, connectivity to lab instruments, and bioinformatics code. Pre-competitive collaboration can extend even further to how AI models themselves are built. Federated learning allows companies to update a shared global model without sharing their underlying datasets with competitors. This approach has already shown significant improvement in AI models for small molecules, and can likely have an even larger impact for large molecules if companies invest in it together.

The era of biotech AI

The era of rational drug design — an atom by atom, computer-aided approach to designing drugs for a specific target — started more than 30 years ago. It had a tremendous impact, leading to breakthroughs

for debilitating diseases like cystic fibrosis. We are now entering an era of AI-driven drug discovery, which promises to be AI's greatest contribution to humanity by finding treatments for the thousands of currently untreatable diseases.

If the bulk of biopharma R&D continues to operate as it does today, a treasure trove of new drugs will be created that may never make their way to patients. Applying AI beyond drug discovery and using it to reinvent all elements of R&D will ensure that doesn't happen.

Ashu Singhal and Sajith Wickramasekara are the co-founders of Benchling, a company that provides a cloud platform for life sciences R&D. Singhal is the company's president; Wickramasekara is its CEO.

STAT

Regulating artificial intelligence doesn't have to be complicated, some experts say

By Brittany Trang | **JULY 03, 2024**

ADOBE

Artificial intelligence has the potential to revolutionize how drugs are discovered and change how hospitals deliver care to patients. But AI also comes with the risk of irreparable harm and perpetuating historic inequities.

Would-be health care AI regulators have been spinning in circles trying to figure out how to use AI safely. Industry bodies, investors, Congress, and federal agencies are unable to agree on which voluntary AI validation frameworks will help ensure that patients are safe. These questions have pitted lawmakers against the FDA and venture capitalists against the Coalition for Health AI (CHAI) and its Big Tech partners.

The National Academies on Tuesday zoomed out, discussing how to manage AI risk across all industries. At the event — one in a series of workshops building on the National Institute of Standards and Technology (NIST)'s AI Risk Management Framework — speakers largely rejected the notion that AI is a beast so different from other technologies that it needs totally new approaches.

Experts highlighted strategies the government already uses to regulate the drug, transportation, and financial industries that could and are being used to regulate AI. They also suggested self-regulating market approaches, like AI companies insuring themselves against adverse outcomes from their products, as a way to force companies to exercise caution.

"When people say things like, 'There's no American law that regulates AI,' what I say is, 'Could you name a law that doesn't regulate AI?'" said Erie Meyer, chief technologist and senior advisor at the Consumer Financial Protection Bureau. "There's no universe in which AI is off limits to the civil rights that Americans fought and in some cases died for."

Meyer gave the example of a chatbot study her office did last year. The

resulting report reminded financial institutions that regardless of the technology used to power the chatbots, the decision tree or AI was still subject to the same financial laws as their human bank tellers.

The same principle extends to laws about being able to explain why someone's credit was denied, said Meyer, where the CFPB rejects any "black box" technology — AI or not — that spits out a decision no one can explain. Just last week, lawmakers sent Medicare a letter urging them to do the same for algorithms that result in care denials for Medicare Advantage patients.

Industry insiders have said that regulating AI is difficult or inadvisable because it's moving too fast and that laws could impede innovation or make using the AI tools cumbersome. But Meyer pointed out how the Fair Credit Reporting Act, which is over 50 years old, was designed to grow with changes in technology. It remains relevant today because it focuses on inputs, not just hard-to-predict outputs, she said.

"In the Fair Credit Reporting Act, there are permissible purposes [for using consumers' data] and there are impermissible purposes," she said. "One thing I will highlight that is not on the list of permissible purposes is training a model." The agency is highlighting this point in a forthcoming rule in an effort to stem any businesses that are pivoting to brokering their customers' data, she said.

Workshop co-chair Ben Shneiderman, a professor emeritus of computer science at the University of Maryland, said that when it comes to AI, he couldn't believe how quickly Microsoft abandoned its system of rigorously testing new features with different groups and eventually broadening features as they worked out bugs. "That 40-year

success story has been tossed aside…I don't understand where the aberrant belief [comes from] that the only way to test it was the phrase, 'let's put it out publicly and let's see what happens,'" he said. Is there a way to get companies to voluntarily roll out AI products in a way that raises the level of safety?, he asked.

Jill Crisman, executive director of safety testing lab UL's Digital Safety Research Institute, proposed a bold idea for AI companies to self-regulate how much risk they are willing to take with their products. "What if AI providers, or anyone who's deploying an AI system, actually provided reparations for any public harm that their AI system enabled?" she said. She admitted it might not be possible, but that "this would do more to establish trust in a company's products than anything else than I can actually think of."

Marc Rotenberg, the founder and executive director of the Center for AI and Digital Policy, an independent nonprofit research organization, pointed out that when building a building, "the building inspector [must come] to certify that the construction has been done properly, [and] then and only then can insurance companies become the source of reparations." Would the reparation payments be paid out by insurance companies, or voluntarily?

Insurance companies only take on risk when the risks are local and they can spread out the payouts on a broad risk pool, said Crisman. "What happens when you have an insurance system where the ramifications could be global? This is why I think it's the companies themselves who are collecting the profits from these that should be in self-insuring and providing reparations," she said.

Simply having to do the risk calculations and figure out the monetary value of the damage the AI could cause would make developers more cautious about the benefit of rolling out the product widely or quickly versus its costs. And conversely, a warranty that an AI model not only meets minimum standards but exceeds them would be a powerful marketing tool, she said.

But Crisman still said that regulation of AI would help rein in risks in a very resistant industry, especially one that's trying to stay ahead of other countries like China. The automotive industry fought seatbelts for a long time because making their product safer cost them money, she said.

The best argument for standards not impeding technological progress, said Shneiderman, is that a new Boeing or Airbus plane takes $1 billion and a year to get FAA certification, and drug trials also take years. "When we're dealing with life-critical and consequential applications, should we expect a systematic process?" he said.

"I will say that the automotive industry actually did seatbelts because it was regulated and it was required," said Crisman. "And so this is why I say that regulation helps."

STAT

As Washington mulls new AI rules, startup leaders want a seat at the table

By Mohana Ravindranath | **JUNE 05, 2024**

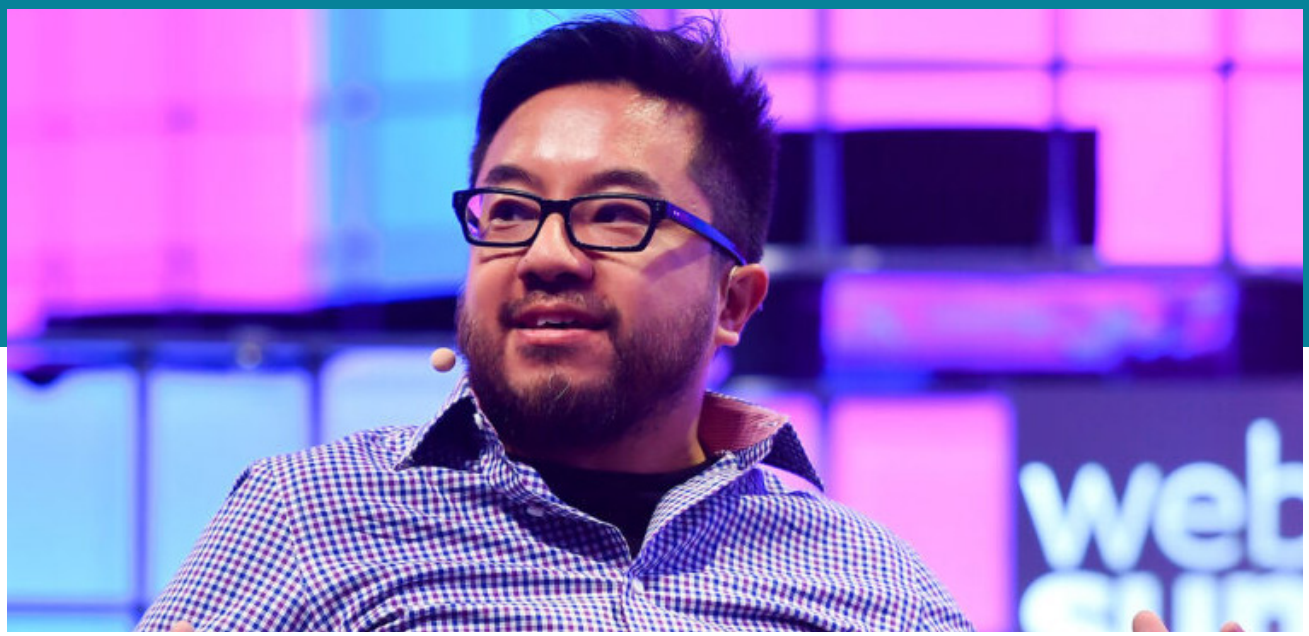

Garry Tan, CEO of Y Combinator
SEB DALY/WEB SUMMI

Washington — In San Francisco, Y Combinator firebrand chief executive Garry Tan is one of the startup economy's most recognizable figures. Well before he became the leader of the coveted startup accelerator that bred global brands like Airbnb, Stripe and OpenAI, Tan was a self-made founder who sold his own company to Twitter.

But here in the nation's political capital, Tan is far less known. He's hoping to grow his influence here, offering up his vast network of founders as a source of expertise for regulators struggling to wrap their heads around new technologies like artificial intelligence and how best to rein it in.

"We're happy to be a phone-a-friend for people who are trying to understand tech," he told the Economic Club of Washington recently when asked how he'd measure success during Y Combinator's first-ever policy-focused visit. "'Little tech' represents competition, and consumer choice," he said later at a separate event in a dimly lit Senate room filled with a few dozen Congressional staffers.

As Congress and the White House weigh AI's risks to consumers and patients against its potential, Tan is just the latest business leader to wave the "little tech" flag in Washington and offer its wealth of technical knowledge and the promise of a dynamic, vibrant startup economy that's only achievable if proposed regulations don't hamstring startups.

The new urgency among startup leaders to be heard in D.C. comes as federal regulators work against the clock to put in place guardrails to ensure the breakneck pace of development and deployment of artificial intelligence in health care doesn't harm patients. Venture investors and startups worry that the new rules are likely to favor big tech because well-resourced corporations like Microsoft and Google can spend more time and resources to make their case in D.C.

Venture capital firm Andreessen Horowitz, which registered to lobby last year on issues such as Medicare and AI, is also championing the

thousands of small companies developing AI and software tools for sectors like health, finance, and retail, health general partner Julie Yoo told STAT. The firm, an early backer of now global companies like Airbnb and Coinbase, has urged regulators not to burden startups with arbitrary hurdles that only well-resourced, big tech companies can afford to clear. Andreessen Horowitz has a large portfolio of prominent biology and health startups, including Omada, Komodo, Turquoise, Devoted Health and Benchling.

The Consumer Technology Association, which hosts CES, the sprawling Las Vegas gadget showcase, has also convened a handful of educational briefings on AI and health tech this year along with the Digital Health Caucus, a bipartisan group of lawmakers led by Reps. Troy Balderson (R-Ohio) and Robin Kelly (D-Ill).

"We define 'success' as seeing more members on Capitol Hill understanding and asking questions about health-focused technology and the opportunities to improve care," CTA digital health vice president René Quashie told STAT last month. The lobbying group is fielding growing interest from startups eager to shape AI policy, a spokesperson said.

There's still concern among tech leaders that Washington doesn't have the expertise to regulate AI, whether in Congress or the executive branch. A new proposal from the White House aims to address that, mandating agencies appoint their own chief AI officers.

Health tech investor Bob Kocher, a former Obama administration official who helped write the Affordable Care Act, said he's met recently with the Senate, the Food and Drug Administration and the Centers for

Medicare and Medicaid Services, but he didn't think they're currently equipped to regulate AI in a way that nurtures new ideas but also protects patients.

"I do think [members of Congress] have a humility, and that they have a desire to not impose regulations right now on AI because they don't know how to do it in a way that protects patients, that rewards innovation, that's going to work right," Kocher, a partner at venture capital firm Venrock, said at STAT's Breakthrough Summit West in May. Venrock has backed health tech companies such as Lyra and Included Health.

"Congress couldn't possibly have the expertise to do this today," he said, noting that it would have to hire much larger staff with sector-specific expertise. "We need more participation in government — we need to have people who bring expertise, to bring it to government," Kocher said, adding, "I think government needs the help."

Regulators aren't necessarily on the same page yet. While large swaths of a White House executive order on responsible AI use are "probably on the right track overall," Tan said, he told Federal Trade Commission head Lina Khan late last year that a measure proposing that models above a certain computational size be subject to regulation was "widely mocked," and "sparked backlash from our community, especially AI startups."

"The EO could have also benefited from less explicit definitions" including that computing threshold, Y Combinator's policy head Luther Lowe told STAT. In Congress, a new "roadmap" authored by a bipartisan Senate AI working group outlines broad policy goals,

but doesn't do enough to support small companies competing with behemoths, he said. Y Combinator is also growing its health portfolio after hiring its first health group partner in 2021.

"We need flexible frameworks to address this fast-moving space, and being too prescriptive could hinder the ability for startups to be nimble as they leverage AI to build something people want," Lowe said. "[I]gnoring the perspective and interests of startups by writing policy that kills them before they ever have a chance to get off the ground would be one of the most damning mistakes lawmakers can make."

Transcarent, a health benefits management tech company led by industry veteran Glen Tullman, is among startups escalating their D.C. influence. The company, which rolled out an AI care navigation tool in May, joined CTA last year; a Transcarent senior product manager who's also a doctor testified during a House Energy & Commerce hearing on health AI late last year.

Government affairs lead Leslie Krigstein, herself a health tech policy veteran, told STAT that government leaders have been largely "candid about the knowledge gap between the health care issues they've come to know…and what AI means."

In the early stages of AI policy-making, Krigstein said Transcarent and CTA have focused more on educating regulators than on pushing specific policy measures.

"We've just tried to kind of be in the room and talk about what our use cases are, and what we'd like to see: Congress taking the stepwise informed approach," she said, adding, "for those of us who don't have

compliance departments that are tens or hundreds of people deep, the burden's going to be greater on innovators."