United States House of Representatives

Committee on Energy & Commerce

Subcommittee on Health

Hearing on

"Understanding How AI Is Changing Health Care"

November 29, 2023

Dr. Benjamin Nguyen, MD

Transcarent

Chairman Guthrie, Ranking Member Eshoo, Chairwoman McMorris Rodgers, Ranking Member Pallone, and distinguished members of the Committee, it is my pleasure to appear before you today to discuss how artificial intelligence is changing healthcare.

My name is Dr. Benjamin Nguyen. I am a Senior Product Manager at Transcarent, leading our AI team, which is tasked with improving the care experience, the quality of care, and the affordability of care by expanding Transcarent's affiliated virtual clinic's suite of AI tools while maintaining the highest standards for patient safety. I am a medical doctor, completed my degree at the Keck Medical School at the University of Southern California, and I have worked at the intersection of technology and care delivery throughout my career, with a special focus on artificial intelligence.

About Transcarent

Transcarent was founded to make it easy for people to access high quality, affordable care and offer greater choice and control over healthcare for those that pay for care, including healthcare consumers (our Members) and employer-sponsored group health plans. Transcarent is not a standalone health plan. Rather, services offered through Transcarent make the Member's major medical plan easier to understand and use.

Transcarent cuts through the complexity of the current healthcare system, making it easier for people to access high-quality, affordable care. Our platform is personalized for each Member, offers access to physicians and an on-demand care team, and a connected ecosystem of high-quality, in-person care and virtual point solutions, Transcarent guides Members confidently to the right level of care.

Transcarent Members have access to care through digital guidance along with the clinicians in Transcarent's affiliated virtual Clinic as well as from high-quality providers in their local communities. We provide one place for all their health and care needs. Transcarent Members access complex care resources for behavioral health, musculoskeletal or oncology care virtually or in-person, often in value-based care arrangements that align incentives between those providing care and those paying for care. We use the power of technology to scale access to high-quality care, regardless of a person's geography, income, education, ethnicity, disabilities, gender, or language.

Today I will speak to the applications of AI within our affiliated virtual clinic. At Transcarent, we believe AI holds great promise to revolutionize care delivery for the better – to speed access to treatment, reduce administrative burdens on clinicians, and democratize access to healthcare resources (it no longer matters where you live if you want to talk with an endocrinologist, whether in a rural area or a major metropolitan city.) We also believe that AI should be used only for the right reasons – it should not be used to deny access to appropriate, medically necessary care or in a way that diminishes the quality or safety of that care.

Scope of This Statement

AI is an extremely broad term, used inconsistently even by those who work in the industry. There are many valuable types of AI, such as decision support systems or voice recognition AI systems. I chose to focus this statement on a specific, new kind of AI known as Generative AI--the underlying technology powering products like ChatGPT. I chose to focus on Generative AI systems because they have advanced at an unimaginable velocity in recent years and will have a significant impact on the healthcare industry.

AI Technology Has Changed

To understand the importance of the recommendations made in this statement, it is important to understand recent changes in AI technology. It is common to think about the advancement of technology along a smooth continuum, like a train track going up a hill. But sometimes, there are technological leaps that are so great that they propel us rapidly forward into a different world. A very large leap occurred in the AI industry over the last few years, specifically in a subtype of Generative AI technology known as generative large language models (LLMs). The magnitude of this change is so large that it is like going from locomotives to powered flight. Like the advancement from locomotives to planes, this leap brings new risks and opportunities. It is crucial that Members of Congress understand what has changed because the nature of the changes will drive the nature of the risks and opportunities in American healthcare.

Generative LLMs can be used in many domains, but they are most well-known for being the engine behind next-generation chatbots, like ChatGPT. As AI engines, Generative LLMs are computing models that can parse and generate fluent written language. Three key characteristics distinguish Generative LLMs from the AI that most people are used to seeing.

First, chatbots that are powered by Generative LLMs are surprisingly good at handling complex topics, understanding nuanced contexts, and weaving them into their decision-making and answers. This means that unlike chatbots built on older AI technology, chatbots built on Generative LLMs can adapt to their users' complex requests and needs, holding complex and nuanced extended conversations, and handling tasks that require logical reasoning. These abilities set these new AI models apart from previous generations, enabling them to engage in tasks that AI systems of the past could not.

Second, Generative LLMs are by nature, very flexible. This contrasts with older "narrow" AI systems. Older, Narrow AI systems are built in such a way that they can only function within a narrow domain. If I build a Narrow AI system that answers basic questions about the branches of the government, I cannot ask it to generate a poem in the style of Shakespeare or to give me career advice unless I rebuild the AI model from scratch. On the other hand, a single modern Generative LLM given sufficient inputs can handle any of these tasks easily, despite not having been explicitly designed for any of them. Some Generative AI products also combine powerful image recognition AI systems with Generative LLMs, enabling them to interpret scanned documents or even images and other media and incorporate the context into their answers.

Third, and perhaps the most important distinguishing factor, is that the full capabilities and limitations of these new AI systems are not fully known, even by their developers. This feature of Generative AI means that bias is even more challenging to detect when compared to older

Narrow AI systems. There are many reasons for this. First, as with most modern AI systems, Generative LLMs must be "trained" on massive volumes of written language – the ultimate compendium of human experience. It therefore inherits the inherent biases of that experience through the data used to train the model. Second, during their development phase, most Generative LLMs depend heavily on a large number of human workers who "grade" their performance, helping to "teach" the AI system to hone its capabilities. Bias can be introduced in this stage inadvertently, both through bias introduced by the selection of human workers and through the subconscious biases of these workers themselves. Third, Generative LLMs, especially when used in chatbots, can accept open-ended written inputs and create open-ended written outputs. Compare this to a simple Narrow AI system for predicting whether it rains or not. We know that the worst this Narrow AI system can do is mispredict the weather because it is a simpler machine with fewer parts that can break. On the other hand, a Generative LLM chatbot can take any written input, and it may respond in countless ways based on its internal probabilities. I might have no idea that it can be biased until I ask it the right questions, using the right combination of words. When the ways that an AI system can be used multiply, so do the potential avenues of harm.

Opportunities in Healthcare

In healthcare, some tasks are difficult because they require deep expertise and clinical judgment, and there are tasks that are difficult because they are tedious and labor-intensive, yet necessary. Generative LLMs will someday be applied to the former. But today, they can already excel at the latter. And that's good news because for patients, administrators, clinicians, and the multitude of supporting staff charged with delivering care to patients, those tedious and labor-intensive tasks are the work that makes our healthcare industry tick. For instance, it has been estimated that 30% of healthcare costs are administrative.[1] With proper design and testing, Generative LLMs are very well suited to the everyday administrative work of information synthesis, documentation creation, and form filling which can reduce administrative burden and staff burnout. These are obvious use cases, and they will be Generative AI's first proving ground in healthcare.

But these uses, as incredibly valuable as they are, will be relatively incremental in terms of impact in comparison to coming applications in medical triage, navigation, and personalization. The great leaps will come from transforming the way that everyday Americans relate to their healthcare system. The American healthcare system is complicated and confusing. To get healthcare, the average patient exerts immense effort. They wade through the complexity of insurance coverage, copays, deductibles, in- and out-of-network clinicians and benefits, doctors' schedules, and office locations – all before getting care. It is a wonder that patients even have the energy left over to make sound, well-informed medical decisions. Generative AI is a major step towards simplifying the healthcare system so patients can focus on their health and the health of their loved ones. A well-designed Generative AI health navigation product would allow a patient to simply state that he needs an appointment with a doctor – and an appointment will be made with an in-network doctor who speaks the language he is most comfortable with, who has a time available during the patient's lunch

[1] High U.S. health care spending: Where is it all going? (2023). *Commonwealth Fund*.

break and is within a 10-minute drive. He can and should only need to focus on getting healthy.

The complexity and dynamics of healthcare also lead to another kind of inequity - "one size fits all." As the country moves to a more value-based care approach, we need to deliver more personalized, high-impact care. However, different patients have different needs. Some want to ask their doctor more questions than others. Some have a harder time understanding the side effects of their medications. Others need help deciphering lifestyle modifications to stave off heart disease. Yet we are all bound by the typical 10–15-minute office visit, and we receive the same post-visit brochures, written with the same talking points. The fact that most Americans rely heavily on the internet to supplement their understanding reflects our failure to meet these diverse needs.[2] A Generative AI product could create educational materials that are tailored to the patient's level of health literacy and education, in the patient's preferred language, including content that addresses the actual questions that the patient asked during their visit, helping to relieve their most salient healthcare concerns. A Generative AI chatbot working hand in hand with a doctor can enable a medical visit to go for as long as it needs to, by conversing with the patient, answering all their questions fluently in the language of their choosing, and allowing them to ask for clarification or simplification. This is one way that Generative AI can reduce the structural biases of the healthcare system – by moving us from "one size fits all" to "many sizes for many needs."

Other impactful uses of Generative AI in healthcare include but are not limited to:

- personalized behavioral modification plans to reduce chronic disease risk, considering specific patient preferences, budget and geographic constraints, and family support;
- assisting patients to apply behavioral health coping strategies in conjunction with medication and talk therapy;
- organizing and synthesizing research and expert information for healthcare practitioners and even patients and families;
- training and education for medical professionals, personalized to individual needs and levels of expertise; and
- guiding and supporting families and patients through complex medical decisions.

How Transcarent Uses AI

When a patient comes to Transcarent's affiliated virtual clinic, an AI assistant immediately begins to gather information from them about the reason for their visit, organizing it for the clinician. By the time the clinician greets the patient, they have a detailed and relevant summary of the patient's symptoms and history. This reduces administrative burden and allows the clinician to spend their time focused on diagnosis, treatment decisions, and working in partnership with the patient on follow-up or preventive care. This approach serves both clinicians and patients.

[2] Finney Rutten LJ, Blake KD, Greenberg-Worisek AJ, Allen SV, Moser RP, Hesse BW. Online Health Information Seeking Among US Adults: Measuring Progress Toward a Healthy People 2020 Objective. Public Health Reports. 2019;134(6):617-625. doi:10.1177/0033354919874074

Perspectives for the Future

My own journey from medical education to building AI products has given me some perspective that I hope will be valuable to the members of this Subcommittee.

First – norms of responsible AI healthcare product design must be established deliberately and thoughtfully. This does not necessarily mean new regulation – but it does mean that stakeholders across the healthcare industry must come together to share knowledge and establish what responsible AI use in healthcare means. The AI industry's biggest players have established their own internal AI safety divisions due to the potential harm that AI can cause. Applying AI in a unique and multifaceted industry like healthcare will create unique and multifaceted ways that it can cause harm. For example, we may want to establish the principle that AI should never be used to make decisions to deny appropriate medically necessary care.

In the absence of healthcare-specific frameworks, we have established a proposed set of core principles to govern our AI product development at Transcarent:

> (1) Patient safety - Patient safety means that any AI products we build must include redundant safety mechanisms as a core part of their design – safety systems surrounding AI are our "table stakes" features, not afterthoughts. We build AI products to augment and enable clinician decision-making and diagnoses. The responsibility to diagnose patients remains with the clinician.
>
> (2) AI equity end-to-end - AI equity end-to-end means that we consider how decisions along the entire product development process will impact equity. Will a data set used in development inject bias into our AI product? Or will applying AI technology in a different domain than was originally intended lead to biased results?
>
> (3) Patient- and clinician-centered AI – We design our AI products to enhance the experience of our patients and clinicians. That means that we consider them key stakeholders during the design process, and work to ensure that our AI products earn their trust and are built to serve their needs first and foremost.

Second, there is a significant and growing gap in AI talent within the healthcare industry's practitioners and leaders. Healthcare's unique challenges and opportunities mean that we need to develop internal expertise in AI and Generative AI. Even amongst AI companies and experts, there is not good agreement yet on *how* to even measure the capabilities and safety risks of new AI technologies, much less how to mitigate the risks. The development of healthcare AI products that serve all Americans equitably *demands* active participation from all levels of the healthcare system, from executives to those providing care to patients at the bedside. Generative AI systems are so complex that the study of their biases, capabilities, and potential for malfeasance or beneficence to society has become a discipline unto itself (known as "alignment research"). The researchers who study Generative AI are often more akin to biologists running experiments than engineers building machines.

Thus, we must also study Generative AI as a discipline within healthcare, or we risk leaping into the future without understanding the consequences unique to our industry. Having gone to medical school, I know that the training of individuals in a discipline orthogonal to the hard medical sciences does not come naturally to our institutions. Most medical students and

doctors do not have pathways, nor do medical institutions have the funding or expertise to create them. We need incentives, frameworks, and funding to create these pathways if we want to be able to ensure Generative AI can achieve its potential to revolutionize our healthcare system.

In conclusion, I believe that the integration of Generative AI into healthcare, as with many other industries, holds great promise. At the same time, it is critical to balance the promise of AI with safeguards that will be necessary to build and preserve trust across all stakeholders.

Thank you for your attention and I am available to address any questions you may have.