

Goldman Sachs Exchanges

Generative AI: Hype or truly transformative?

Sarah Guo, Founder, Conviction

**Gary Marcus, Professor Emeritus of Psychology and Neural
Science,**

New York University

Kash Rangan, U.S. software equity analyst,

Goldman Sachs Research

Eric Sheridan, U.S. internet equity research analyst,

Goldman Sachs Research

Allison Nathan, senior strategist, Goldman Sachs Research

Dates of recording

June 8, 21 and 27; July 11

Allison Nathan: Since the November release of ChatGPT, a breakthrough generative artificial intelligence tool, investor interest in generative AI technology has surged. Technology stocks have substantially outperformed, driving the sizable equity rally in the first half of the year. But is this technology truly transformative or overhyped? I'm Allison Nathan and this is Goldman Sachs Exchanges.

On this special episode, we're breaking down the hype around generative AI that was the topic of our most recent Top of Mind report, now available on GS.com. We dig into the disruptive potential of generative AI technology and

whether it warrants the current investor enthusiasm. We speak with Sarah Guo, founder of AI-focused venture capital firm Conviction; Gary Marcus, professor emeritus of psychology and neuroscience at New York University; and Goldman Sachs's equity research analysts Kash Rangan and Eric Sheridan, who cover the US software and Internet sectors.

We first asked Rangan and Guo about what differentiates generative AI technology from previous developments in AI. Rangan explains that the technology's ability to create new content using natural language prompts are its two key transformative features. What is it about this technology that is so exciting?

Kash Rangan: What is different about generative AI is it can generate content and generate code, and content can be text content. It can be video content. It can be image content. And we have not been able to do that before using natural language commands. That intersection of using natural language commands, English to generate these types of content we're talking about has never been done before, and that is what is generative about generative AI, as opposed to AI which is training a computer to make simulations, predictions of human behavior, business outcomes, that sort of thing.

Allison Nathan: Guo further explains that, unlike previous iterations of AI technology, generative AI doesn't require humans to write code to perform specific tasks or collect large amounts of training data. The foundational models that lie at the heart of generative AI have already been trained on large datasets to handle a wide variety of tasks with natural language prompts. So generative AI capabilities are now widely accessible, which, Guo says, is ushering in an era of Software 3.0.

There's been so much focus on generative AI. What's

differentiating about it that is getting so much focus?

Sarah Guo: I do think it's breaking change. It's not just more machine learning. The practical impact is you have these new capabilities that are much more general and more powerful and enable very different product user experiences. In traditional machine learning, you're trying to do one task at a time, and that might be text-to-speech. It might be image recognition. It might be summarization. But it's often very expensive to go through the machine learning development cycle of data collection and engineering to get something to work at a certain quality level against a given task. And that's one of the reasons we haven't seen this very widespread adoption.

But what has happened in the era of foundation models that we're entering now is you just have many more capabilities that are out-of-the-box available by an API and in the open-source as well. And I think one useful framing in terms of what that increased accessibility means is there is a technologist named Andrej Karpathy. He led the auto-pilot team at Tesla. And in 2017, he coined this term “Software 2.0.”

And Software 1.0 you can think of as software that is deterministic code written by humans, function by foundation, very expensive labor. Software 2.0 you can think of as optimizing a network to make decisions for you where the main work is not actually writing the software but it's collecting training data to train a neural network. Classical machine learning.

I think we're really entering this era of Software 3.0, where you can do a lot more out of the box with foundation models. You don't need to collect the training data yourself for many use cases. And so it suddenly becomes much less expensive and the reasoning capabilities of these models that have so much general knowledge is now available to any company that wants to go invest in it.

And I think the biggest opportunity for investors is understanding the basis of competition and the margin structure for many industries I think will change. And it's not just limited to things that were purely software. A lot of things that more traditionally services I think are now within the realm of being served by more software.

Allison Nathan: Can you just give an example of that

kind of opportunity?

Sarah Guo: Sure, we're investors in a company in the legal space. And what the company competes with is not let's say, like, a traditional software that might be workflow oriented, serving the legal industry, but actually it has the ability to do the legal work that was too expensive to do prior. So imagine you want to understand 25,000 contracts at once. At the quality of your junior legal team doing that first pass or very expensive due diligence work that wouldn't have been done before, that's services and you're enabling legal firms to do higher level work for less.

Another example of something that would not have traditionally felt like a software space would be an image generation company. And I think if you'd asked venture capitalists, public market investors, they would not have described illustration and image generation as an important or large software market today but it suddenly is. There's a significant number of companies that look like that.

I'll give one more example. We're investors in this company in the analytics automation space. They don't replace any

particular part of the modern data stack today, but they replace the work of low-level data analysts serving your business functions, answering questions from your data warehouse. And so I think automating that low-level of knowledge work, which is not a software function today, which is done by humans so that you can go do higher level analytical work, is another example of -- that's not an existing software market; it's scope expansion.

Allison Nathan: But even if AI technology ultimately proves transformative, has the hype around what the technology can actually deliver and what the market is pricing gone too far at this point? Gary Marcus, who's done extensive research on AI, cautions that generative AI technology isn't yet as transformative as it's often made out to be.

Gary Marcus: At the core of all of these tools is basically something like auto-complete that is trained on a very substantial fragment of the Internet. So almost everything that you see is some application of a tool that is pretty good at auto-complete. It's not pretty good at understanding the world, and so it makes stuff up. And it turns out that these tools are really good at some things like computer

programming, where you're writing code. A lot of it is predictable, and so the system auto-completes for you so that's a good application.

A slightly disconcerting application is if people use the same tools for medicine. Auto-complete really isn't sophisticated enough for what you want. People try to make it sound like these machines reason like people, and they don't. They're really not that much like humans. At most, they do some of what Daniel Kahneman might call System One, automatic statistical analyst. They do very little of what Kahneman would call System Two of deliberate reasoning. They just don't do that well.

They do learn things, but most of what they learn is about the statistics about words. And now you have a second thing called reinforcement learning with human feedback that's learning in certain contexts what's socially appropriate to say when. So there's definitely some learning in these systems, but they're not learning abstract ideas. They're not learning abstract concepts. There's a lot of learning that we do about the world that they don't really do.

Like, I'm looking at you, and I can see something about the room that you're in and something about your emotional expressions. I could make some guesses about you. I'm trying to understand the world that you're in, and these systems just don't do that. They don't have a curiosity about the world. They don't have a representation of the world.

Allison Nathan: Marcus warns that the intelligence of AI systems is overhyped and argues that we're nowhere near achieving artificial general intelligence.

Gary Marcus: Artificial intelligence tools really are having a material impact on our life right now. These tools, for example, can be used to create misinformation and that's going to affect the 2024 election pretty significantly. So there's some real positive impact right now on computer programmers and some real negative impact on elections, but the intelligence of these systems is really not that great.

There was a recent case study where somebody claimed that the system passed the undergraduate exams in engineering and computer science at MIT, and people got

very excited about this. And it turned out that the methodology was terrible. People love narratives about the current machines are changing everything, they're going to displace all the works. I mean, it's very exciting to talk about all of this stuff; it's not really true. Yes, they are having an impact now. They really are changing society now; that is true. But they're not that smart. There's still a long way to go before even we solve driverless cars that we've been talking about for a decade.

And other people are worrying what if robots take over the world; they seem so smart? They're not really that smart. They're nowhere near taking over the world. I used to joke, if the robots come for you, first thing you do is close the door. I made that joke four years ago, and it's still true. Robots still can't open doors. So, yes, there's real-world impact. No, these things are not really anywhere near artificial general intelligence. There's a lot of people who think that AGI is imminent. I think they're mostly wrong.

I feel like we're in the age of alchemy where people kind of knew what was going on but hadn't really figured out chemistry and didn't really have laws for what they were doing. Didn't really know what experiments would work

and what wouldn't. But they had their finger on something happening, and we're kind of like that now. People can smell what AI might look like now, but we haven't really developed a rich enough, sophisticated enough understanding about how to build intelligence in machines. There's no reason to think we won't get there eventually.

I'm often cast as the pessimist, but I had a debate with Grady Booch, a very well-known software architect, and I took the optimistic position that we would get to artificial general intelligence in this century. And he took the less optimistic position that it might take even longer than that. I think we'll get there eventually. There are some people who, from a philosophical standpoint, would say intelligence simply isn't something you can build in machines, but I've never been particularly impressed by that kind of argument.

You might make arguments like a machine will never feel pain. That might be true. But I think arguing that a machine would never understand what a person is feeling when they're in pain and what they might do as a consequence of feeling that pain -- take a painkiller or go to a doctor, whatever -- that's a totally different thing. I think

that machines will eventually develop a much clearer understanding of human beings. They'll be much more reliable. They won't make stuff up as much. I think we'll solve all of that. It might take another five years. More likely, it'll take another 20 or 40 or 50.

Allison Nathan: And when it comes to markets, even Guo, who's optimistic about the technology's transformative potential, sees areas of hype today. Although she says she's less focused on valuation and more focused on identifying AI-led companies that could become important.

Sarah Guo: A very just common pattern in investing at any stage is people misjudging the time scale of these different technology shifts. Overall, I've made a very fundamental bet that this is a decade-plus transition that is going to be very important and drive real value shift and value creation.

But in the near term, at least what's happened in much of the private market early investing ecosystem, is there's general technologist agreement that this is a very important shift. And you have shoots of real data that get

investors very excited. And so you have a huge cadre of investors who are trying to figure out how to leverage the trend or figure out what their risk profile is around it.

And one of the things that happens is, while they're still developing a deeper understanding of a space, it's very easy to anchor to more visible heuristics. So if you're an early investor, one of the heuristics might be, "Did the person work at OpenAI or DeepMind as a researcher?" That is an easier question to answer than, "Is this particular product or research thesis going to work?"

And don't get me wrong. These are amazing research organizations, but I'd say where people have obvious profiles and the theses match a pattern that has existed in the past as a valuable software category, their popular idea right now is vector databases. I've probably seen 25 vector database companies. Then there's a lot of investor enthusiasm. And where there's limited supply and a lot of enthusiasm, prices go up and diligence is less strong. So I'd say there are areas of hype.

Allison Nathan: In this more challenging valuation environment, how do you think about investing? And what

gives you pause?

Sarah Guo: The great thing about being an early stage investor is the simplification for us is can we imagine a company going from zero to one into 100 of revenue? Can it be an important company? And our business is actually a little less focused on particular valuation within a range, right? We want to be in the best people in the companies that matter as an early stage venture firm.

Allison Nathan: And Goldman Sachs's US Internet analyst Eric Sheridan is comforted that the vast majority of companies that have recently outperformed on the AI theme are still trading at relatively reasonable multiples. Here he is talking about this.

Eric Sheridan: You never know if you're in a bubble or not. Of the companies leading this that have outperformed the market over the last four to six months, the vast majority of them are still trading at relatively reasonable multiples to gap enrages per share. Bubbles are typically about enterprise value to eyeballs or clicks or addressable market dynamics being talked about and sharing euphoria as a driver of valuation as opposed to

what the right multiple on net income is to pay. So that's very different than other bubbles that we've all lived through.

Allison Nathan: Rangan agrees that this time feels different. He points out that, unlike other large technology cycles in the past, the most powerful technology companies in the world are driving this shift. So he argues AI probably isn't in a hype cycle.

As much excitement that there is around the potential of the technology, is there some sense at this point that it's overblown? We have seen moments around Web 3.0, the Metaverse that have fizzled out. So why won't this be a similar scenario.

Kash Rangan: Generally, big technology cycles are originated by upstarts. We had the birth of distributed computing back in the early '90s, advocated by a little company called Oracle against IBM, which was all about the mainframe cycle. It took a while for that distributed technology to become more pervasive because it was one big voice, IBM, that said, "This is the established way of doing things." So generally, switching over from old cycles

to new cycles, there are obstacles and established business practices, established technologies that get in the way of adopting a new way of doing things.

Given the cloud cycle, going from distributed computing to cloud, the existing on-prem systems and technology providers said the cloud is not safe, it doesn't scale that well, and it's not that economical, etc. And you had to overcome years of objection before the cloud really finally got some footing. And then the larger established companies, when they adopted cloud, there was harmony of voices telling the buyer this is okay; you can do this.

This cycle, what is so different is you don't have this being led by upstarts. It is being driven by some of the most powerful technology companies on the planet. When there is a unanimous verdict among the technology providers that this is actually happening, this is real, and when customers start to get interested, it's not hype. We're having real-world discussions with CIOs of big decision makers, and the productivity benefits have already been tallied up and they look absolutely, stunningly impressive. So it doesn't feel like a hype cycle.

Allison Nathan: Finally, we dig into the risks around investing in AI technology right now. Guo warns that discerning between AI marketing and AI reality among the current enthusiasm could prove difficult, as could regulatory backlash.

Sarah Guo: Distinguishing between AI marketing and AI reality is going to be hard work for investors. I do think it's amazing the founder and executive commitment from public companies to the AI trend so quickly. But painting everything with AI marketing in your earnings calls isn't going to do much good if it doesn't translate into revenue, cost improvement, increased scope of business, resourcing AI. And the inventor's dilemma of you having a manual services business and you need to figure out what to do from a labor perspective when you can suddenly automate a ton of it is a very complicated question for people to answer. And I do think that's a risk for many public companies.

Allison Nathan: What's your advice to investors on how to separate out those who are just marketing and those who are actually putting in the resources?

Sarah Guo: I look at outcomes. If I think about the outcome of your consumer business and you shipped, let's say, chat-based features or semantics search or something new that involves AI in the interface, how do you already measure your business? It's in engagement or in transactions or in ad inventory. And if the AI is valuable, then it should show up in one of those metrics.

And don't get me wrong, there's delay in product development with all of these things, but I tend to think of it as are we driving engagement or revenue or changing the margin structure? And how will we know? As early stage investors, we're investing in the thesis, so we're giving companies time to go show those results. But we tend to be pretty clear about are you moving the needle on one of these metrics that's important to the business already? Or are you enabling us to do entirely net new work? And if you are, do customers like it? Are we selling that product? Is it a new and better experience?

Another significant risk is this strong backlash in regulation and concerns about risk around AI. I think there are very real practical concerns around abuse of these technologies. Bias. Disinformation. Cyber security.

If you have very general tools, just like the Internet, people are going to use them for good things and for bad things, and you have to invest in mitigations as well.

But I think we're so early in some of the impacts that are going to be very important for society, like in science and education, in healthcare, that it would be a real shame if we ended up regulating the industry to a halt.

Allison Nathan: Meanwhile, Sheridan is closely monitoring how consumer behavior might evolve. For example, he warns that, if consumers shift away from the traditional search engine, that could upend existing businesses.

Eric Sheridan: From a consumer Internet standpoint, the potential for changed consumption or computing habits is pretty much at the forefront of our industry analysis. And I've lived through multiple cycles of a certain thing is going to disrupt the search engine. Mobile was supposed to do it. Social media was supposed to do it. Even Amazon versus Google was a debate investors had five, six years ago. We'll just all go to Amazon and type what we want into the search box there

and that'll be end of search.

So these are big open-ended debates about how computing habits might change, but there are much more variable impacts down the stack of what that might mean for the way business models are constructed to generate demand, convert demand, and then match it with supply that we're going to have to continue to monitor and watch for.

You have whole industries built on certain elements on aggregated supply and generated demand. And if the consumer behavior changes away from the search engine or the search engine has to change as a result of it, you could see wildly different economic outcomes than what investors are used to today.

Allison Nathan: And Rangan worries that the more pervasive the technology becomes, the less valuable it might be. Here he is again.

Kash Rangan: The one risk that I think about is generative AI becomes too accessible. Today, because of the expertise that you need to train large language models on, you need to have some supervision, very few computer scientists are

specialized in generative AI. It's a brand new field. But still, as to how quickly these elements are going to be learning from the data and content, if it ends up being shockingly quick and we have widespread diffusion of what we consider to be very special. So if it's not special, then where is the premium pricing for this? And if everybody has it, then we're spending a bunch of CapEx to effectively level the playing field. And so where do we get paid for all this? What is the incremental value?

Unless it really impacts productivity so positively that you're willing to pay extra for it. And if you're not willing to pay extra for it or if it gets competitive so quickly that you can use somebody else's generative AI technology on top of your application, you can substitute the generative AI with half a dozen things that are available that can do the same thing. Where's the differentiation? That worries me.

Allison Nathan: With questions about the transformative potential of generative AI continuing to swirl, we'll be keeping close tabs on this technology and its impacts. I'll leave it there for now.

If you enjoyed this show, we hope you follow on your

platform of choice and tune in next week for another episode of Goldman Sachs Exchanges.

Make sure to like, share, and leave a comment on Apple Podcasts, Spotify, Google, or wherever you listen to your podcasts. And if you'd like to learn more, visit [GS.com](https://www.gs.com) and sign up for Briefings, a weekly newsletter from Goldman Sachs about trends shaping markets, industries, and the global economy.

Speaker: All price references and market forecasts correspond to the date of this recording. This podcast should not be copied, distributed, published, or reproduced in whole or in part. The information contained in this podcast does not constitute research or a recommendation from any Goldman Sachs entity to the listener. Neither Goldman Sachs nor any of its affiliates makes any representation or warranty as to the accuracy or completeness of the statements or any information contained in this podcast and any liability therefore; including in respect of direct, indirect, or consequential loss or damage is expressly disclaimed. The views expressed in this podcast are not necessarily those of Goldman Sachs, and Goldman Sachs is not providing any

financial, economic, legal, accounting, or tax advice or recommendations in this podcast. In addition, the receipt of this podcast by any listener is not to be taken as constituting the giving of investment advice by Goldman Sachs to that listener nor to constitute such person a client of any Goldman Sachs entity.

This podcast should not be copied, distributed, published or reproduced, in whole or in part. The information contained in this recording was obtained from publicly available sources, has not been independently verified by Goldman Sachs, may not be current, and Goldman Sachs has no obligation to provide any updates or changes. All price references and market forecasts are as of the date of recording. This podcast is not a product of Goldman Sachs Global Investment Research and the information contained in this podcast is not financial research. The views and opinions expressed in this podcast are not necessarily those of Goldman Sachs and may differ from the views and opinions of other departments or divisions of Goldman Sachs and its affiliates. Goldman Sachs is not providing any financial, economic, legal, accounting, or tax advice or recommendations in this podcast. The information contained in this podcast does not constitute investment advice or an offer to buy or sell securities from any Goldman Sachs entity to the listener and should not be relied upon to evaluate any potential transaction. In addition, the receipt of this podcast by any listener is not to be taken to constitute such person a client of any Goldman Sachs entity. Neither Goldman Sachs nor any of its affiliates makes any representation or warranty, express or implied, as to the accuracy or completeness of the statements or any information contained in this podcast and any liability therefore (including in respect of direct, indirect or consequential loss or damage) is expressly disclaimed. This transcript should not be copied, distributed, published, or reproduced, in whole or in part, or disclosed by any recipient to any other person. The information contained in

this transcript does not constitute a recommendation from any Goldman Sachs entity to the recipient. Neither Goldman Sachs nor any of its affiliates makes any representation or warranty, express or implied, as to the accuracy or completeness of the statements or any information contained in this transcript and any liability therefor (including in respect of direct, indirect, or consequential loss or damage) are expressly disclaimed. The views expressed in this transcript are not necessarily those of Goldman Sachs, and Goldman Sachs is not providing any financial, economic, legal, accounting, or tax advice or recommendations in this transcript. In addition, the receipt of this transcript by any recipient is not to be taken as constituting the giving of investment advice by Goldman Sachs to that recipient, nor to constitute such person a client of any Goldman Sachs entity. This transcript is provided in conjunction with the associated video/audio content for convenience. The content of this transcript may differ from the associated video/audio, please consult the original content as the definitive source. Goldman Sachs is not responsible for any errors in the transcript.