Talks at GS

Transforming Work Productivity with AI: Glean CEO Arvind Jain

Date of recording: October 15, 2025

Arvind Jain: Imagine an AI that actually is a personal companion to you, that walks with you wherever you go, listens to every word that you speak, listens to every word that you hear, is in all the meetings with you. And now it actually comes to you proactively and does some of that work for you.

Kim Posnett: Hello.

Arvind Jain: Hi, there.

Kim Posnett: So you actually got the idea for Glean while you were running Rubrik. So you observed that scattered knowledge, quote/unquote, was a problem plaguing most organizations. Tell us about Glean's origin story, and what was the void you were trying to fill?

Arvind Jain: Yeah, well, first of all, thank you for having me. Really excited to be here. The origins of Glean, I

mean, as you mentioned, I was doing a startup before Glean, Rubrik. And we had good success in that company, grew it rapidly. So within five years, we were more than 2,000 employees.

And while the business was growing very rapidly, before that but as we scaled and we got to so many employees, productivity sort of tanked in the company. We got into this place where engineers, we had tripled the size of the engineering team, but yet we couldn't actually make software releases anymore. We used to take about 3-6 months to make one, and we're sitting with one release for more than 12 months.

Our sales teams were not selling as much product as they used to be able to. So there was a big problem in the company, and we tried to investigate. You know, talk to people, see what's wrong. And in our pulse surveys, the first, the biggest complaint that our employees raised was they could not find anything in this company. "When I need information, I don't know where to go look for it. We have 300 different cloud-based systems. Our knowledge is spread across all of these systems. I also don't know who to go and ask for help when I need help because there's no

employee directory. There's no concept of who works in what."

So that was the problem that we had, and when I saw that that was the first coding question on our pulse survey, I said, "Well, I mean, this is an easy problem. I'm familiar with this. I've worked on search before. And, yes, I also find it difficult to find information inside the company, so let's go buy a product that will connect all of our enterprise information and make it searchable." And we discovered that there was nothing to buy.

And so that sort of started a thought in my mind that maybe I should go back to search and this time solve it for people in their work lives.

Kim Posnett: So tell us a bit about Glean's capabilities today, and what can you do now that was sort of unheard of 5-6 years ago?

Arvind Jain: Yeah, so when we started out, our simple was mission. We wanted to build a Google-like product inside your business. Bring Glean inside your company, connect it to all of your enterprise systems, and give

employees one place where they can go and ask questions. And we'll surface the right information from wherever it is within your company and do it in a safe and secure way. Don't leak information. Like, know who's asking for information and only show them things that they have access to. So those are the things we're trying to solve. Very simple. Like, you know, just wanted to build a Google.

At the time, in early 2019, transformers were not as big of a thing as they are now, but they were there. And people in the search industry, you know, we knew it quite well, like what transformers do and what impact it was going to make on search. In fact, the origins of all of this, all of these new networks and transformers is actually search. At Google, we started these projects to make Google Search better.

So we had familiarity with that, so we took that as one of the core technologies inside our company, which actually made us the first enterprise company in the world. We didn't really have any clue, like, you know, how fast the technology was going to evolve. It was really very good for us in terms of conceptually understanding information so

we could build a good search product.

But then over the years, the product has evolved. So instead of looking like Google inside a company, now you can think of it as a version of ChatGPT that knows everything about your company, your business, how work happens inside your company, who are the different people, what do they do. So you can actually come into Glean and you can ask any question, Glean will go and work on that, use the latest model technologies but also all of that enterprise context and information to answer those questions back for you. So it sort of looks and feels like ChatGPT, but think of it as a superset of ChatGPT because of all of the enterprise context.

And then over the years, we've also evolved to take this platform that we've built, which connects all of your enterprise information and data, bring it in one place. We make this platform available to agent builders. So if you are trying to take a business process and agentify that with AI, you'll be able to do it with ease with the Glean AI platform.

Kim Posnett: So you've said publicly that the majority

of work we do today will not be done by us humans and inside of five years.

Arvind Jain: Yeah.

Kim Posnett: You've also said this is about agentic AI. You've also said every employee should have a team of AI agents to help them get their work done. So how does this all play out? And what does productivity look like for all of us in the future?

Arvind Jain: Yeah. I mean, see, today AI is actually already much better at doing a lot of tasks that we have to do as humans because AI, for one, doesn't get tired. It has infinite stamina. It can go and read as many documents, can spend a lot of time researching, analyzing, and that capability is something that we don't have.

I'll give an example. Let's say I'm a product manager trying to figure out what are the new features I should be building in our product. As a human, I could go and maybe talk to 10 or 15 of our most important customers and go and ask them, "Well, tell us the top five things you want us to work on." And then I can take that set and come up with

another top ten things that I should prioritize for my business.

But if you think about a large business with thousands of customers, you have a lot more information that you are constantly hearing from your customers. A lot of calls that have happened. A lot of email exchanges. A lot of customer support tickets. And there are these hundreds of thousands of artifacts that AI can now actually go in and then look for what features customers have requested over the last one year, organize it thematically, and actually come up with an idea for what your road map should be based on that more comprehensive view of what all of your customers want.

So that's an example of things that AI can do so much better than humans today already. But in terms of how our work as knowledge workers, how it's going to evolve with AI in the future, you can take any role, whether you're an engineer, a customer support person, somebody in sales business, investment professional, majority of your time today is spent doing things that are not particularly strategic in nature. There's a lot of legwork that you have to do to actually try to get to that place where you can then

go and make strategic decisions. And a lot of that work will be something that AI will be able to assist with in the future.

But I think the real cool thing that we think is going to happen is AI, like, today it's very reactive. So today, if you want to actually use AI in your work, well, first, you have to be motivated to use it. Second, you need to know what AI can do for you. And both of those things are hard. Even in our busy work lives, people don't have time to explore what AI can do for them. And they also sort of don't really have that imagination. In fact, nobody knows what AI can do for us tried. Most of the capabilities of these models is uncharted territory. It has to be explored by humans over the next few years to know what AI can do for them.

So we sort of end up in this situation where, today, AI is not making as much impact in businesses, in our work lives. Usage is actually fairly low. And if we change, if you sort of flip the model and actually bring AI to people as opposed to you going to ChatGPT or you going to Glean and asking it questions while giving it tasks, imagine an AI that actually is a personal companion to you, that walks with you wherever you go, listens to every word that you speak,

listens to every word that you hear, is in all the meetings with you, is in systems where you work like email or Slack or documents that you write or read. And with all of that context about you and your work life and tasks that you need to work on, now it actually comes to you proactively and does some of that work for you.

So that is what has to happen first before this concept of AI doing the majority of our work for us is going to happen.

Kim Posnett: By the way, for the entrepreneurs and CEOs in the room who know that they want to integrate AI more into their businesses but don't know necessarily how, what questions to ask, how to drive business transformation, what's your advice to them?

Arvind Jain: Well, I mean, first, I think you just have to have this assumption that AI can do everything for you, and you have to just try it out. Give it more and more complex tasks. And I think for those who tried AI three months back, four months back and were disappointed, that's sort of the real challenge because if you sort of give up based on that then you're going to miss out on all the great improvements that are coming now.

In the last six months, AI models have fundamentally changed in their capability, in their about to do a lot more complex set of tasks than what they could do before. So I think the number one thing that we can all do is just keep using it and change our instinct.

Like, one of the things that I do now is that whenever I have questions -- like, before in our limited release, my instinct would be to go call my CTO or my sales leader and ask them the question. But slowly, I started to realize that takes too long. Half the time they don't pick up the phone, and they don't actually have the answer themselves.

Kim Posnett: Need to work on that.

Arvind Jain: They have to go and ask somebody else to get it. And so now my instinct is that whenever I have any questions, I first call and ask Glean. Glean is connected to all of our enterprise context. It knows about every conversation between any two people. It has read every single document. And then it has the power of these models that we use from OpenAI and Anthropic and Google.

And so a lot of these complex questions that I have, I just first go and ask them in Glean and let it come back with an answer. And maybe it's wrong. Maybe AI hallucinated, but it doesn't matter. That's a good start for me. And then I just share that session sometimes with my CTO and say that, "Hey, I had this technical question. Is it the right answer? Tell me." If it is then I'll just use it.

But in that process, the whole thing took me two or three minutes as opposed to .

Kim Posnett: Hours and days.

Arvind Jain: Or a week because it has to trickle down all the way through the organization to actually get the right response. So that is actually the one thing. Changing our instinct and actually using it more and more, you get to then also learn how to prompt, how to really effectively leverage AI.

Kim Posnett: You just mentioned hallucinations. So a big concern certainly across enterprise usage is hallucinations. Glean, you seem to be able to avoid

hallucinations. How do you do that?

Arvind Jain: Well, I mean, I think, first, the models are actually getting better at dealing with hallucinations or making things up because with the thinking models, they actually now have been trained to doubt their own work. Before, you could actually ask the model, "Hey, prove that one plus one is equal to three," and it'll actually go and prove it because it won't actually really have that instinct to say that, "Hey, look, I may be wrong in the work that I'm doing." So that's a new capability that models have now.

So you can design them in a way where you ask, for example, you ask a question to a model. It uses the knowledge inside the company or out in the world and comes back with an answer. But then it actually sort of says, "Look, I may be wrong. Maybe I should actually get a little more evidence. There's some doubt that I have in what information I've seen so far. Maybe the information is out of date."

And so that's sort of what is happening now with the models. They think more. They think like us. And it's the same thing, if you imagine your CEO comes and asks you

for a question on some topic. So then you will actually be super careful in terms of getting the right answer to them. You won't actually just do a search, get a quick answer, and give it back to them or talk to one person, get an answer, and then share it immediately with your CEO. You'll probably do some more. As a human, you know that this is important. "I should actually double check, make sure that the answer is absolutely right." So you'll actually do a little bit more sort of fact seeking and evidence verification. And the thinking models now do the same.

So hallucinations, in general, I would say is not a problem in enterprise anymore. The bigger problem in enterprise is actually: Can you even get models the right knowledge to work on so that they can do their work properly? If you give it stale information to work on, the answers are going to be stale. And so what we do is, in Glean, that is the work that we do is understanding the enterprise context well, understanding what information is up to date and written by subject matter experts so that you can trust it more than others. And so that's the work that we do in terms of assembling the right input for the models to work on.

And then the way that actually our product works is that when people come and ask questions and we answer them, we don't just show them the answer from AI. Line by line, we'll actually also try to point them to that human-generated information that was used to produce that response. So that way people can sort of actually build that habit of sort of fact checking while reading their answers.

Kim Posnett: So you've said that the agentic AI that you launched last year called Glean Agents is on pace to support one billion agent actions by year end. What does that look like?

Arvind Jain: Well, I mean, agents are being used quite actively now by our customers. The way our product works is that it's horizontal in nature. You first deploy Glean in your company, and you connect it with all of your enterprise systems. And once you do that, now you can actually go and build agents across different functions.

For example, maybe you want to automate requests that your HR team gets from people. Maybe you want to actually automate customer service ticketing. You want to

automate redlining of contracts that your legal teams have to go through. So these are all different business processes. And you can actually go and build agents to automate all of these business processes.

And so what we're seeing is that now enterprises are actually going quite broad in terms of where they want to bring AI inside their company. No function wants to be left behind. There's significant pressure from CEOs to see value from AI across all of their different business functions, so we're definitely seeing this infinite uptick in agent usage across different use.

Kim Posnett: So let's shift to competition. So in your sector specifically, there is fierce competition when it comes to AI for the enterprise, including everyone from your former employer, Google, to OpenAI to Dropbox, Salesforce, on and on. How do you view your capabilities as different from the others? And how do you think about competing with that crowd?

Arvind Jain: You know, when we started the company in 2019, nobody wanted to touch enterprise search. In fact, I had investors who told me that, "Hey, here's a blank check.

Write the amount that you want to raise." And since I was an entrepreneur already, I had success behind me, so people wanted to invest in me. But the person who said that, you know, gave me a blank check and says, "You know, as long as you don't work on enterprise search, go fill this out." The space was like that. There was no success in the field. There was no large company that came out of it. And everybody thought that search inside a company is sort of like a vitamin. It's not a painkiller. Not a good business to build.

But for me, it was something that, like, I felt the pain and I saw it in my own startup. And at Rubrik we saw this was the biggest problem that people were complaining about. And so I had conviction that this was an important problem to solve. So we went ahead and did it.

And for the first four years of our company, we were alone. We were the best product in the domain, but there was no other product so it was easy to be best. And we built a customer base slowly and steadily. We had to evangelize a lot. We had to tell customers that investing in your people, making it easy for them to find information means something. Go and invest in it. So that was sort of the

mode that we were in.

And then the ChatGPT moment happened, and the light bulb went on. People started to think about, "Well, if I had something like ChatGPT inside my company and I could actually train a model that could take all of my company's data and information and then just work like ChatGPT where I can ask questions and I can get the right answers back, it would be so amazing." And it was interesting that that's exactly what we were at Glean. So that sort of helped us really scale in the market. We got that demand that we were waiting for when our product actually was really ready and mature.

But now, search has become a core foundation for making AI work in the enterprise. Any agent that you want to build inside an enterprise ultimately needs information inside your business to work on, to do its work. And typically, the way you get to that information is through a search system. So the search technology became very foundational for making AI work in the enterprise, and every large software company saw that. And that's why they're all coming. It's not just Google or OpenAI. It's also Microsoft and Salesforce and ServiceNow, Workday,

Atlassian, you name it. Everybody's coming into this market.

So the question for us is then: How do we distinguish ourselves from them? I think we're ahead of course because we started many, many years before the rest of the industry came into this area. So today our product works the best, and we have to work hard to actually maintain that lead. Ultimately, that's the only answer always, which is you have to move fast. You have to stay focused, which is easy for us because this is the only business we have versus all these other companies. They have a lot of .

Kim Posnett: You've talked publicly about not overestimating competition. Maybe tell the audience what you mean by that and how it might apply to their businesses and sectors.

Arvind Jain: Well, typically, in our case at least, the demand for AI, like, what businesses want to do with AI is so high that even if all these companies that we talked about, if we all came together and we created swim lanes and started to not compete with each other, we still cannot deliver to the demands that the market has. So there's

plenty of customers to work with.

And that's why I always feel like competition really doesn't matter. More often than not, you actually kill yourself than somebody else coming and killing you. It's like building on innovation is hard. Once the teams get large, keeping them aligned, keeping them moving together in the same direction, keeping everyone motivated, executing fast, these are the real challenges. That's sort of what we're always focused on.

If you execute fast, you'll do well because it's very, very hard to execute fast and especially hard for the larger companies.

Kim Posnett: Okay, so let's talk about that. Let's talk about the challenges that you faced as a founder from this very quick rise of launching Glean in 2019 as a startup to growing to an AI unicorn times multiple with over 800 employees. And I think you said this in a recent interview, that the AI industry doesn't allow anyone room to breathe. And I think you were talking about that crushing pace of innovation and growth.

Arvind Jain: Yeah.

Kim Posnett: You've said your mission is to expand human potential to do extraordinary work. Some might argue -- by the way, I am not in this camp -- but some might argue that the proliferation of AI agents and assistants paradoxically could shrink human potential by diminishing a worker's critical thinking, problem-solving skills. How do you answer that question? And how do you make sure that Glean is actually augmenting human intelligence and not replacing it or diminishing it?

Arvind Jain: Yeah. It's the same thing as when you think about the calculator. There was a time when people could do multiplication of two numbers in their mind, and now you can't. And so the --

Kim Posnett: Long division.

Arvind Jain: Or long division, yes. But did that shrink human potential or expand it? I think we would agree that I think we're probably better, we can do more things than what we could do before. And that one thing that we used to be able to do we don't have to because I got the

calculator in my pocket, right?

So I think it's the same thing with AI. First, I think we're overestimating. I actually strictly feel like AI is an augmentation tool for us. It's just going to replace some tasks, but it's not replacing any human as such anywhere in my opinion. We will create the world's largest enterprises, and we've brought a lot of significant productivity, even some gains to them. But if you go to any individual manager within that company, they're not actually saying that, "Hey, look, I'm going to actually shrink my team by half." And they don't want to. And actually in some sense that is right because I think your business is going to be proportional to your employee base. I fundamentally believe in that concept.

Of course, you have to be efficient. As efficient as your competitors or more. So I think with AI, I feel it just takes us one or two levels up in terms of the kind of work that we do. In fact, it actually makes room for us to do more strategic thinking.

I'll give this example. So I was going to meet the head of AI at AWS last week. And I'd been sort of running around,

didn't really have time to prepare for that meeting. And so the night before, I actually go in Glean and I ask this question that, "Hey, I have this meeting. Just help me out. You know, what should I do?" And it does a lot of that heavy lifting, builds some contexts. It says that there are some new product launches that they've announced. Maybe that's a good thing to talk about. You know, here are some potential integration ideas for how our product can work with them. And it's sort of doing that work that I would have my chief-of-staff do for me who would then ask a few other people. But I sort of now start with this. I get a really great starting point, but I'm still thinking about it and thinking strategically with that good base to start with.

So while a lot of work was done for me, I feel like I was able to spend more time not struggling with information gathering and getting that sort of work that I don't like to do. But I was able to sort of think and spend time on more strategic work with that. I don't know. I feel pretty optimistic that AI is actually going to really make more time for us to do more creative things.

Kim Posnett: I agree. So prior to founding two of your companies, you spent over a decade at Google, and you

were working as one of the company's first engineers and eventually ran Teams, Search, Maps, YouTube. There's an interesting story about the time Larry Page and the executive team asked you to make the Internet run faster. So tell us about that conversation, and what did it teach you?

Arvind Jain: Yeah, so there are a few interesting ones there. So first, this started with -- so my role at Google was at some point I became the speed czar. The role of the speed czar was to make sure that all Google products work super fast. Like, they load in less than one second. People don't have to wait.

And so I used to work on it, and it was actually I was probably one of the most disliked persons in the company because I was in the way of all new great features that would make our products a little bit slower. So there was this one meeting when I was with Larry, and I update on, "Hey, look, we made this product a little bit faster." At that point, I made 400 milliseconds faster. And, like, he says, "Like, well, what about Hotmail?" I think he was trying to use Hotmail at the time or something. Or CNN. I forgot what he was talking about.

And I said, "What's that got to do with Google?" And he says, "But no, I still go there and I don't want to waste time." And then he says, "You know, we should make all of the websites on the Web faster. We should make them more secure." So for me, it was hard to sort of comprehend. But then I asked him and he says, "Well, we're the world's largest Internet company. If the Internet is faster and better, we're going to make more money. We're going to make more business because everybody's going to spend more time on it."

Anyway, that was the first conversation. But then he asked me to come back with a plan for how we're going to make the Internet fast now. And so I worked on it for a month or so and then came back and shared a plan. And it was utterly underwhelming for him. He looked at it and I think he didn't even look up in most of the presentation that I was making. And then his final comment after we were done was that, "I want you to go and come back with a new plan. And in that plan, make sure that you show me how you are actually going to spend $10 billion." And that's what he said.

I thought that maybe the meeting ran out of time and he didn't complete the other part that we also have to make $100 billion back from it. But no, that was not on his mind. He just wanted me to think big. He just wanted to make sure that, when we work on projects, that we don't sort of constrain ourselves by how much money we can invest in these projects or can we make these things profitable or not. First, go and see all the things that you can do in the world to make, in this case, make the Internet fast. And then we'll figure out what we will and what we won't.

And that was a learning for me. It's sort of like, of course, people don't change. And I'm still sort of constrained with my own way of thinking, but I always try to remember that lesson when we build our startups that sort of go and think big.

Kim Posnett: Thank you, Arvind.

Arvind Jain: Ah, thank you.

not necessarily reflect the institutional views of Goldman Sachs or its affiliates. The material provided is intended for informational purposes only, and does not constitute investment advice, a recommendation from any Goldman Sachs entity to take any particular action, or an offer or solicitation to purchase or sell any securities or financial products. This material may contain forward-looking statements. Past performance is not indicative of future results. Neither Goldman Sachs nor any of its affiliates make any representations or warranties, express or implied, as to the accuracy or completeness of the statements or information contained herein and disclaim any liability whatsoever for reliance on such information for any purpose. Each name of a third-party organization mentioned is the property of the company to which it relates, is used here strictly for informational and identification purposes only and is not used to imply any ownership or license rights between any such company and Goldman Sachs.

A transcript is provided for convenience and may differ from the original video or audio content. Goldman Sachs is not responsible for any errors in the transcript. This material should not be copied, distributed, published, or

reproduced in whole or in part or disclosed by any recipient to any other person without the express written consent of Goldman Sachs.