



SuperDataScience

**SDS PODCAST  
EPISODE 981:  
HOW DATA  
ENGINEERS ARE  
“10X’ING”  
THEMSELVES WITH  
AGENTS, FEAT. MATT  
GLICKMAN**



Jon Krohn: 00:00:00 Whether you're aware of it or not, February 2026 was a huge moment in time, an event horizon, as my guest today describes it, where everything changed for computing, for AI, and for society. Welcome to another episode of the SuperDataScience Podcast. My guest today is Matt Glickman, who spent nearly 25 years at Goldman Sachs before jumping to Snowflake when he sensed big opportunity there. Now, he senses big opportunity in AI, game-changing opportunity with his new startup, Genesis Computing. They are automating data engineering with agents and doing it more accurately at a scale that humans wouldn't ever be able to do. Hear about this and lots of other society changing developments in today's episode. This episode of Super Data Science is made possible by Anthropic, Acceldata, and Cisco.

00:00:52 Matt, welcome to the SuperDataScience Podcast, a treat to have you here. How you doing today?

Matt Glickman: 00:00:56 I'm doing great. It's a beautiful day in New York.

Jon Krohn: 00:00:58 It is.

Matt Glickman: 00:00:59 No more snow.

Jon Krohn: 00:01:00 Exactly.

Matt Glickman: 00:01:00 At least not until next week.

Jon Krohn: 00:01:02 Yeah. It looks like it might actually snow next week. Literally. Which is crazy. But yeah, at this time, we're recording in mid-March in New York City and we have the most beautiful day of the year so far. I personally can't wait to get out of the studio, but before that, we're going to have a great time in the studio with Matt Glickman here. And it's going to be a fascinating episode because you've taken the plunge and you're doing something that seems like such a huge opportunity.



- Matt Glickman: 00:01:32 For the second time. It is like insanity. You don't repeat yourself or maybe repeat yourself if you're insane. Maybe it's a little bit of that.
- Jon Krohn: 00:01:41 Yeah. Yeah. Well, you got to keep going. So you spent two decades at Goldman Sachs.
- Matt Glickman: 00:01:45 Just missed a 25-year anniversary. I know because at Goldman, there is an anniversary dinner at the 25th year and I just missed it. I got the clock at 20, which is a symbolic.
- Jon Krohn: 00:01:57 Kind of
- Matt Glickman: 00:01:57 Gesture of maybe you should be ... Time is ticking. And then fate had other plans.
- Jon Krohn: 00:02:06 Tell us about that.
- Matt Glickman: 00:02:07 Yeah. So again, basically I joined Goldman straight out of school, very fortunate, and grew with the company, grew with the teams I was a part of, and basically was running the data platform team for the quants, which was a big part of how Goldman was successful. And financial crisis hit and we were fortunate in that we had a platform that everyone was using to basically manage risk and were able to add onto that platform data tools and data platforms that allowed these quants to be able to help Goldman manage the crisis. And we saw how powerful that was where the power of a platform and the power of being able to give these people who understood the business and understood the tech, the tools to do really data at scale, navigate the crisis and that was great, but then we'd basically given them a taste of what was possible.
- 00:03:08 And it soon became something where instead of just something that you could do these kind of emergency analysis to want to run your entire business on this



consolidated platform. This was pre-cloud, pre-Snowflake, pre-everything.

- Jon Krohn: 00:03:23 So their own servers. Own servers,
- Matt Glickman: 00:03:25 Physical servers. Physical servers, legacy kind of tech. But giving you a hint of what was possible, but we realized that while everyone started talking about the big data problem, what we had was a big user problem. We just ideally wanted to run the entire firm off of the same copy of the data. And then you could put more data and put more elements of Goldman's business into this one place and everyone could operate on it,
- 00:03:51 Which obviously became the bottleneck. And I remember making a promise to the head of the quant team saying, "Promise me, we're not going to have the entire firm running off of this one database." And of course, I'm like, "Well, that'd be crazy." Of course, I couldn't stop it because ... And this actually happened again when I went to the asset management side of the business and same thing happened. We showed what was possible. Everyone wanted to run the entire business on it. It couldn't scale. Fortunately, one of my colleagues knew one of the founders of Snowflake at the time. And this was, again, early days, 2013, maybe early 2014. And they came to pitch Snowflake to a room full of Goldman Skeptics. Fortunately, I was forwarded an invite. I show up the last person, last seat, and I sit down, happen to sit down next to Snowflake's founder, who the first and last time I've seen him in a tie, Benoit, brilliant, brilliant system designer and architect.
- 00:04:51 And he basically laid out the solution, which was basically if you decouple compute from storage and you leverage the power of elasticity in the cloud, you can solve this big user problem. And again, 2014, and it was early. And I remember there was one guy in the room who was trying to figure this out with me, and he says, "Well, this



is great and all to Benoit, but we're in Goldman Sachs. It's 2014. We're not going to the cloud by this craziness." And I remember, and I've made fun of them since then, but the other snowflakes in the room, the sales team, were not breathing. They were thinking like, "Maybe we're going to close Goldman Sachs in this early days." And then, but Benmov, of course, laid it out plainly.

00:05:40 By the time Snowflake is ready to be on- prem, Goldman will be in the cloud. And everyone kind of laughed inside and the Snowflake people cried, but he was 100% correct because that was the answer. You never were going to be able to keep track of that scale and be able to adapt without a architecture built without a problem and the scale of what the cloud could give you, particularly for this big user problem. So I left that meeting and I figured, well, I can either pretend that meeting didn't happen or try to somehow take those learnings and apply them into this limited on- prem capability hat, or I could reach out and see if I can help.

Jon Krohn: 00:06:24 No kidding.

Matt Glickman: 00:06:25 And reached out.

Jon Krohn: 00:06:26 Wow.

Matt Glickman: 00:06:27 And I hadn't reached out for a job since I had joined, so I didn't even know

Jon Krohn: 00:06:30 What to

Matt Glickman: 00:06:30 Do. And there was literally one, this was, again, very early days. There was one job listed on the snowflake.net. It wasn't even. Com yet.net website for marketing. I knew nothing about marketing, but I'm like, I'll just apply and maybe ...

00:06:47 I found out later they thought this was how Goldman evaluated its vendors by applying for job, but I reached



out. Next thing you knew, I was talking to then CEO, Bob Logley, the following weekend, and it basically laid out my understanding of what problem they were trying to solve and how this could not only be applicable for a Goldman in its internal data problems, but really once you're in the cloud, you can effectively connect each of these enterprises together because industry like finance is all interconnected. Data's not being invented in Goldman. It's coming in from data vendors, coming in from markets, and it's coming in and also going between all these players. If you're operating in the cloud, it actually not only gives you that scale, but it gives you the opportunity to interoperate more efficiently by not having data moving around at traditional methods.

00:07:37 But I also understood what it would take to basically get a company like Goldman to adopt this kind of technology. And yeah, so as I've described, and similar to how we've started Genesis, you don't sort of ... I don't think, maybe someone does, but you don't come up a list and like, "I'm going to come up with a list of things I want to do and start a company." For me, and I've seen this with others, opportunities present themselves and you

Jon Krohn: 00:08:04 Can either

Matt Glickman: 00:08:06 Ignore them or you could realize the opportunity and basically kind of run with it. And that's what we did. And then the rest is history.

Jon Krohn: 00:08:17 Snowflake wise.

Matt Glickman: 00:08:18 Snowflake

Jon Krohn: 00:08:19 Wise.

Matt Glickman: 00:08:20 So I went, left Goldman, joined Snowflake, led product there for the first few years. We were talking about in the setup here, I actually ended up spending the first three years by Coastal. So a week in the Bay Area, working with



engineering, a week in New York working with customers.  
It was not-

- Jon Krohn: 00:08:43 Just alternating back and forth.
- Matt Glickman: 00:08:45 I literally took the same JetBlue flight every other week.  
Oh
- Jon Krohn: 00:08:49 My goodness. Did you have an apartment there as well?
- Matt Glickman: 00:08:53 I did. Yeah. So I was basically going to ... And it was surreal because time actually goes by faster and my kids were growing up faster because I miss every other week. But I actually had my up in the air George Clooney moment where taking the same flight literally every other week and one flight, the guy recognizes me and I'm lining up to get to my seat and they make a point of telling you how many miles you have because I was accumulating a crazy amount of miles. And he comes up to me and he's like, "Mr. Glickman, give me a bear hug." We really appreciate the business. I remember people behind me saying, "Who the hell is that guy?" But yeah, but it was very fortunate because as I was describing, New York is still very unique in the enterprise space of just the cross industry concentration.
- 00:09:43 Yeah,
- Jon Krohn: 00:09:43 I think this is a really interesting point to dig into it. It wasn't really something that I had planned for the episode, but I do think this is interesting. And it's important to provide me personally with the confirmation bias that I need to be living in New York- That I can help. And so it's interesting. So I constantly, by being in New York, having my own AI company, hosting an AI podcast, yes, we do shoot episodes in the Bay Area in person, or we do remote recorded episodes very frequently with Bay Area guests, as my listeners have often heard. But anytime I'm visiting the city, I feel like I'm missing out



because the energy around what's happening in AI, the number of meetups that are there, the free drinks, the events everywhere, for people working in AI, and yet you think the best place in the world to start an AI business is New York.

- Matt Glickman: 00:10:35 No question. For enterprise specifically, I'll make this clarification. If you're building something for the consumer space where your customers are all the people in their homes, then the Bay makes more sense because of the engineering concentration. Though I would argue that because of that concentration, trying to retain people is tough out there, but for enterprise, there is no place. I've thought about this a lot. The amount of cross-industry concentration there is in New York doing finance, healthcare, media is unprecedented. And particularly past COVID, being in person is game-changing. We always do Zoom and meet in some teams all day long, but there's still nothing like it. And having an actual office, your team can meet at, you can bring clients to, but then go on site and not have to get on a plane.
- 00:11:36 It's a no-brainer. And just even the concentration of then the operators who are here putting AI to work. And strangely, in my earlier story, finance was the late adopter to the cloud, finance is the early adopter for AI in a big way. And I think that just because of how much operational complexity there is that there really was no answer until now, now a lot of these companies, and same thing for healthcare, also a late adopter is an early adopter for AI, and it's being able to ride that wave. It's actually interesting. The early adopters for cloud was media and gaming, and they're the late adopters now because of the fear of content leakage
- Jon Krohn: 00:12:31 Into the models. That is interesting.



- Matt Glickman: 00:12:34 For finance and healthcare, which is, sure, they have IP in their process, but it's not like they're worried about their banking report templates leaking into models, but also just also the amount of written word that these models are being trained on about these industries is unprecedented. Each frontier model is more and more knowledgeable about this space and about- About finance
- Jon Krohn: 00:13:02 And about care. About finance,
- Matt Glickman: 00:13:04 Healthcare, and that combined with the explosion of power in its decoding ability just presents out of the box models that with the right framework and harnesses and guardrails are accomplishing amazing things. And we're all going to be talking about the February moment, hopefully not in a Cyberdyne kind of way, but that was the moment where everybody basically just realized that this is not slowing down, and it had taken a massive step
- Jon Krohn: 00:13:36 Forward. Yeah. And so you're talking about the capabilities of the models that came out in February and doubling or tripling the length of a human task that it could handle, especially for computer science and machine learning kinds of tasks.
- Matt Glickman: 00:13:48 Yeah. But particularly for us, once again, we're focusing on the data space, data engineering, which are very ... They're very tantalizingly easy, yet very hard, complex workflows that require precision, that require planning and thought and a lot of context. So it's like I was saying, the models understand the industry, they understand the semantics, they understand the business processes. What they're missing is how that is being applied inside an organization.
- 00:14:23 A JP Morgan does something similar, but different than a Goldman Sachs than a city, then a so on. And how that actually gets materialized in people and processes and databases and data flows and all of that, we're hearing a



lot of that discussion on context graph. This is the kind of missing piece. And the other thing that it's typically not even written down. The scary thing, I saw this at Goldman, I saw this, it was always the eliminating factor for clients they're working at Snowflake, there's just not enough reason or time to ever document things. I mean, there's companies are still to this day trying to document everything, create semantic models and all this kind of ... At the end of the day, someone leaves, which they typically do at a high attrition rate, particularly in these high stress kind of data roles, and they leave and knowledge walks out the door.

00:15:23 I was actually there for one of our early design partners, private equity firm. I was there on the day that our data engineering champion was his last day. He never smiled until that day because it's a really rough job, like complex pipeline, not enough people, not enough time. No one knows you exist until it fails. There's a great Simpsons meme that goes around. It's like the Ralph Wiggins guy who was basically like, "I'm a data engineer and no one knows my name." Perfect. Anyway, on that day he was skipping because he's like, "I'm done. I'm burnt. I'm out. " But the guy who was inheriting it looked like he hadn't slept in like 47 hours, bloodshot, hair, rye, trying to, as best as he can, absorb by osmosis, what's going to happen next Thursday when that weird feed goes awry. And that's the big problem is that it's not even just going in and reading documents.

00:16:22 It's not written down.

00:16:23 So part of what's become really interesting is that in order to do data engineering well, you have to actually have the system go and figure things out. Yes, there's going to be human oversight and human kind of steering and guidance, but really what you want to do is what you'd hope a great employee would do would be come in, read everything, look at every database, look at every line of



code, look at every email, look at every communication and glue it together in a full-on context graph, not only defining what things are today, but then figuring out how did we get here? What was the email that kicked off this discussion to change this logic of how we compute our customer retention ratio, which is lost to time. Some of that you can't recover, but a lot of it you can. And again, humans do this.

00:17:23 They read the spreadsheet, they read the email, they look at the data, they ask questions, they then memorialize those questions. All of this is ... I don't think anyone would disagree that that's the holy grail. The interesting part that we've uncovered is it's not something that we do just as a separate thing. It's just part of our Genesis onboarding is to map the entire universe and have that be the starting place to then be productive. And then as a side effect, capture this all because now if that person leaves, I'm saying, it's not going to be ... You still want to have a going away party and miss them, but it's not like knowledge is going to walk out the door.

Jon Krohn: 00:18:11 Right. So yeah, you've been alluding to it, maybe not quite directly, but to solve the problem that you've been describing, you've been describing the problem as well as the solution that Genesis Computing, a company that you founded two years ago is solving. And so you're on a mission to turn enterprises into what you call AI first companies, which is interesting because they're already kind of existing as enterprises, but yield reverse engineer them into becoming AI first companies and so that they can run faster, leaner and smarter by starting every project asking the question, why can't an agent do this work? And I think that February moment is key to being able to ask it that way in the negative.

Matt Glickman: 00:18:51 Yeah. No, you have to force this. I was actually talking to another one of our customers who is actually a hedge fund that just started, and I think it helped them



crystallize that they have a unique advantage versus their legacy competitors. They can start in a way today with no legacy and basically ask this question of saying like, "Well, before we hire anyone else, let's figure out what can we do with just the AIs and then can we ride this wave, particularly for a company like that, that is every dollar that's spent not investments is a draw on their returns." But you have to ask it this way because otherwise you have built up all these processes that were successful up until this moment that we've entered this February moment, and I know it's been talked about, it really does feel like we've entered this event horizon.

00:19:55 There is no turning back and there'll be companies that will embrace it, which as I was describing, surprisingly, the big financials are embracing it more than you'd expect. And it's going to be bifurcation. The ones who embraced it and figured out how to adapt and had a business model that could sustain this transformation will survive and others will literally disappear and it's happening way faster than we all thought. And at this point, we started, my co-founder, Justin Langston and I started Genesis, April will be two years, mainly because of what we were seeing at the time, Snowflake customers trying to embrace this technology to do, which was unblock this bottleneck that all these data teams have. They could never respond fast enough. So they were trying to find a way to use AI to do self-service for their business users and to make for great demos.

00:20:59 And of course, demo would basically plateau and you realize that we get you at 80% and then getting beyond that was almost impossible because you were missing the model capability, but also these harnesses that could really steer these intelligent beings to these solutions. But we realized we were seeing this over and over again and people were trying and failing. And given what we knew about our early view of the technology, because we used this technology to actually win internal hackathons at



Snowflake because we were seeing this early because we were both our hands in engineering, but also in out to customers. And we realized that we could help be this platform that people could start from instead of just starting from the base models. But again, I'd seen the power of platforms at Goldman. I'd seen the power of doing this in a way that would give you scale.

- 00:21:55 And I also saw the power of building something that you're driving revenue for instead of someone internally trying to build for cost. And this has always been the case, and I've seen this of why people adopted Snowflake. It was like you could try to do that and minimize cost, but then you're going to always be beaten out. I remember distinctly the moment I was actually leaving Goldman and going to Snowflake at the time, I had one foot out the door. And so Goldman still had a massive private cloud had similarities to ABS in its concepts, but what would take a minute to provision would take months because it was racking machines. So the equivalent of, for those of you in the audience who've used EC2, and you go there, EC2, you ask for a machine, certain provisions, you hit go and you get it, whatever, a minute and a half later, sometimes faster.
- 00:22:51 So one of the quants I worked with basically filled out the form, asked for the machine, basically three months later, come back and machines are available, gets the provision and it had something in there to, I don't know, to specify what are the mount points. He notices one of them is not necessary. He deletes it figuring it would just apply to his 10 machines. Hits apply, no one had ever done that before. Hits apply takes out half of Goldman because that basically made a global change to the entire
- 00:23:22 Firm. Anyway, so this was, fast forward, the team is blood on the floor trying to clean it up. And I walked over to the guy who was one of the people on the line. I'm like, "It's any constellation, I heard that Amazon, ABS had a



similar issue recently." And he looks up at me with a tier of maybe hope. And he's like, "Really?" I'm like, "Of course not.

- 00:23:46 And I'm like, "And it has nothing to do with engineering pro us." It's the fact of every second that EC2 is not on and billing is money.
- 00:23:58 So that's been optimized, ground down into perfection because it's driving money. With a guy who did this form internally, he was trying to do it as quickly as possible, then move on to other things. It didn't matter if he took a shortcut or not, he did it for as cheap as possible. So this power of doing something for revenue is why when Genesis is solving this problem, so basically Genesis is an Agentic platform that is solely focused on data engineering. We basically realize this is a pain point that almost every enterprise has. There's not enough these people, there's not enough talent. These are hard problems that cross between business understanding and coding and data understanding and it's like all this context. And it's a problem that I've never met. And if someone in the audience is different, feel free to reach out. I've never met a data engineer who wants to do more data engineering.
- 00:24:54 There's just something about it. It's a very painful, underappreciated task that with the right framework around it can be well solved by AI because it loves this stuff. It's a perfect thing where it can bring its general understanding into this space and be applied to go through running every report to make sure it works, looking at every field to make sure it's clean. We have a new customer now who's doing a massive migration and they're working with a traditional consulting company who is basically humans and they have like 30,000 reports that they have to test on this when they migrate everything over. And of course they're not going to run every single report because there's humans. So they'll



pick a sampling and they'll try to see. And then worse, they'll make the business users be the testers because they know what really matters, not all 30,000.

00:26:05 With an AI, you run every single one of those 30,000 and you'll know every single thing that's wrong and it'll be very, very thoughtful and very great documentation on how to cover it.

Jon Krohn: 00:26:16 Yeah. People worry about error rates with agents or with LLMs in general. It seems to me like if you're willing to spend on the tokens to be double checking, triple checking things, you pretty easily, like you talk there about checking over every field, every single item that goes into a report, every line of code, if you're willing to spend on the tokens to double check, triple check everything, you're going to very quickly surpass, especially since this February moment that we've been talking about, you're going to quickly surpass human accuracy levels on, I think basically on a very, very wide range of tasks, including data

Matt Glickman: 00:26:52 Engineering. Yeah, no, definitely. But I think we also, we are given humans way too much credit. Humans make mistakes all the time, and it's just accepted as like, "Well, we're only human." And yet we raised the bar saying, "No, but this other thing that we've created

Jon Krohn: 00:27:07 That somehow is like us." It's a well documented effect.

Matt Glickman: 00:27:09 Right. We just

Jon Krohn: 00:27:10 Expect

Matt Glickman: 00:27:10 It to be better. So number one. Number two is it is something to be concerned about. And I think just saying that I'm going to single shot in YOLO and throw some tokens at this and hope gets the right answer is not the answer either. And part of what we've invested and what we've been building over the last now two years is a



harness that keeps these models on track to not only verify that they get to the right place when they say they are, but actually make them prove it. Part of this next wave is where these models understand how confident they are in their responses. So using that and holding them to task of saying, "Okay, you said you're confident, provide me the actual artifacts that prove that you did these tests, that prove that you ran every one of these 30,000 because for reasons that will somehow maybe or maybe won't understand, the human laziness factor has made it into these models where they will do the minimum necessary to get the answer, right?"

00:28:31 Because you don't want to waste tokens. You don't want to be just going on and on and on if you don't need to. " And if it's like, "Well, I don't really need to run everything of those 30,000 reports, so I'm going to just say it looks good." Versus saying like, "I want proof. I want you to actually give me the results, give me the outputs, and then we'll have another AI that will review it and say like, no, no,

Jon Krohn: 00:28:56 No, no.

Matt Glickman: 00:28:56 There's only 10,000 there." So that kind of checking has been a big part that combined with leveraging this, effectively this internal knowledge graph, this internal context graph, and having that be the two pillars of guiding this intelligent being that we've all created is what's making this work in an enterprise.

Jon Krohn: 00:29:22 It sounds like a known brainer to be taking on these kinds of data engineering agents to help myself and the audience better understand how this works in practice. Are you able to walk us through one or two use cases, maybe anonymized use cases with clients of yours in terms of how you've implemented your solution and what the impact has been on that business?



- Matt Glickman: 00:29:43 Sure. I mean, I guess in the third pillar that I'll kind of mention, and then I'll take you through the example is we're agnostic to the actual data engineering tools and platforms that people are using. These platforms will come and go. The frameworks will come and go. What we're replacing or really augmenting is the people. So now everyone out, to be clear out there, anyone who's panicking, no one's losing their job.
- Jon Krohn: 00:30:15 I have that question coming up.
- Matt Glickman: 00:30:16 No one's losing their job. What's going to happen though is that people are not going to be hired as much. So the next wave of people who would have been the junior data engineer or the junior analyst or the junior operator or just honestly any junior role in enterprises, if you want to get a little more dystopian about it, but the junior level hiring is the problem that AI is basically going to wipe out and it's already starting to happen. Hiring projections are just not going to happen. People have the skills already there. There's no reason to replace them, particularly in the data space because they were limiting the fact they just didn't have enough of them. Now you just can get 10X more power out of them and they can move on to things that they wanted to do and not the tasks that they were basically bogged down to do.
- Jon Krohn: 00:31:03 Yeah. Maybe I'm being too optimistic here, but don't you think there are some scenarios in some organizations where they actually will want to do more hiring because, and even of junior people, because they're being so much more productive with the individuals that they have. They're getting so much, like you said, kind of 10xing the capability there. You were talking about how previously data engineer was this extremely stressful role. They never felt like they had enough time for anything. And I suspect that even if you're 10Xing your capability, there's still a lot to be keeping an eye on and making sure of. And so as that ecosystem grows, you can imagine some



organizations also being like, "It'd be good to have a few more eyes on this.

- Matt Glickman: 00:31:45 " Yeah. No, I think that the few more is the key part of this. My co-founder Justin published a piece that if you haven't seen, it's on our LinkedIn, on Genesis LinkedIn. There's going to be certain people who can actually be these kind of conductors of these AIs at scale and actually use them better than others. He calls it spinning plates. He now has six different agents working in parallel on different things that we're building. I've only gotten to two and a half. So there's going to be certain people who can come in and be these force multipliers and say, maybe that you're working with one agent is doing something with you. I can actually command 10 of them and I'll have the context switching in my head to figure out how to put them all to work. So there'll be some of that hiring.
- 00:32:45 But my opinion, that's going to be kind of the few special hires because the number of people who can do that, it's just a limited kind of audience. I think it's more about just the impact it's going to have. Again, everyone who's doing these kind of operational roles are going to be the ones that now just get a force multiplier applied to them. And I just see it hard pressed to want to bring on more humans when you can just one up 10 agents, 20 agents who do that task instead.
- Jon Krohn: 00:33:27 It wouldn't be surprising to me, and I'm open to your critical feedback on this, but it seems to me like some junior hires might be more likely to be that kind of 10X agent orchestrator as opposed to two and a half X because they could be growing up in this ecosystem where it's like vibe coding first.
- Matt Glickman: 00:33:48 I hope so. The thing I worry about, and this is maybe a little off topic, but I think it's relevant, is that the education system has not yet figured out how to teach AI.



- Jon Krohn: 00:34:05 Yeah, a lot of places haven't. Some places- Some
- Matt Glickman: 00:34:07 Places are.
- Jon Krohn: 00:34:08 And
- Matt Glickman: 00:34:08 I think if that happens, maybe this can change dramatically. But right now, most schools do not in high school and college and don't or discourage people from using AI as part of their process. So it's like a double whammy. Basically, you're going to have companies that are just, if anything, like I said, if anything, they're looking for people who are masters of how to use this tech. Any people come out of school that have learned it despite their schools kind of not encouraging it. So yeah, it's a bit of a double whammy and it's going to happen. I mean, I'm actually trying to give back and work with the schools that my kids have gone to and just to try to express the urgency. And I think there's appetite, but it's tough. I mean, it's, how do you incorporate this? How do you assign papers?
- 00:35:02 What does that mean? The idea that you're going to be able to tell, you're not. That's a farce. No human going to be able to tell and it's just all going to get worse and worse.
- Jon Krohn: 00:35:10 Yeah. Yeah. To bring listeners into kind of an arc of conversations that we've been having about this AI education problem, the human education with AI problem is that a couple of episodes ago in episode 977, we had an NYU professor Kyungung Cho on the show, and he's teaching undergrad machine learning intro course. And for the first time, it's vibe coding first. And he said that it's surprising how many computer science students at NYU have limited to no experience using these tools. And you think those would be the first adopters. But on the note of even younger people, people kind of K to 12, and what are we going to do with education for them? Really



interesting episode 975 a few weeks ago with Zach Cass was exceptional on that. And then we have, I haven't recorded it yet, but I'm expecting that the very next Tuesday episode is going to be with a K to 12 educator who specializes in trying to get some kind of early adoption of these technologies.

00:36:15 That's awesome. Hopefully we'll have some answers for you parents or soon to be parents out there.

Matt Glickman: 00:36:20 I hope so because the dislocation is going to be harsh. Yeah, I hope there is. But back to the question of like, and so again, that socioeconomic problem aside, basically how we engage with customers is basically they onboard Genesis as if they were onboarding new employees on their team. So they connect the system to their platforms, give them credentials to read from their document repository, read from their databases, read from their code repositories. And then they basically define effectively a project or a mission that they want the agents to go on. And the example would be, I want to basically build, I'm an asset manager and I want to basically understand all of my assets by client type, by asset type. I have some raw feeds of things I'm getting from my custodians or my banks and I want to understand everything and create everything I need to actually have an interactive dashboard that I can slice and dice with all my different attributes, which would have been a mass ... Literally, my team built that kind of thing.

00:37:38 It's massive undertaking of just gathering all that data, normalizing it, combining it, linking it, putting all the business logic that typically was hidden away in other applications now into these kind of data flows, producing it as some kind of output. But basically, we've had it where you can literally say like a hand drawn diagram. This is generally what I would want, which actually happens. Typically, if someone understands the business,



they're like, "I want to have these kind of charts and this is like, here, just go and build this and come back to me when you have something." And they come back, whatever, weeks later, months later, and it's half right, half wrong. Basically start there and the system basically starts going. The agents will go and introspect, understand what's already there, understand any kind of, like I said, documentation code or anything as part of that context graph and starts building, building and testing and validating and iterating, taking the best of the coding agent models, but also all this kind of context and using these guardrails that we call blueprints.

00:38:49 So basically we have a set of these, as I'm describing, kind of guardrails that we say, if you're going to extract data, this is the kind of runbook you want to use. You want to extract it, validate it, confirm it, and create a monitor in place to make sure it's always going to be fresh. If you're going to be translating data into some semantic model called a source to target mapping, you're going to basically go field by field and make sure everything ties out and all these kind of things and each step along the way and agents are basically doing that on its own. And the key difference coming back to AI first is that instead of being a copilot where you have to say, "Okay, now do this. Oh wow, that was pretty impressive. Now do that. Oh, you missed it. They should go back and do this.

00:39:35 " We've reversed it instead have the AI is going, working on a task, and when they're not confident or they get stuck, they then come back to the

Jon Krohn: 00:39:45 Human.

Matt Glickman: 00:39:46 And the important thing is when they come back and they say, "No, no, no, no. When I say revenue, this is what I mean by revenue," that then gets memorialized for next time. At the end of the day, but you can now go, we have



these now projects that go on for hours where it's going on and doing things on its own and stopping minimally and particularly as they do this more and more, they can really go all the way through. But it's a combination of that context, these harnesses, these kind of blueprints that they can keep on track. And it's the fact that they can then learn as they go and now become the center of knowledge. And at the end of the day, that's cool and all, but ultimately they solve the problem. These are not assistance. These are great data engineers that are just, you now can scale up on demand.

- Jon Krohn: 00:40:41 Right. Nice. And so your clients have been concentrated primarily in finance healthcare or is there some kind of special- Finance
- Matt Glickman: 00:40:50 Healthcare is dominating.
- 00:40:54 We've been able to attract a surprisingly normally heart onboard set of enterprise customers, mainly because of the way we've chosen to deploy. I have a new appreciation for what Snowflake accomplished back in the day of convincing these enterprises to let their data leave into a SaaS that now had to become a trusted entity. That just doesn't happen. Well, much, definitely not anymore. So the way we deploy, which is definitely the harder way to do it, we deploy Genesis into their environment. In traditional world, you basically give them software and they have to install it and manage it. With the magic of AI, we're basically giving them something to install with an AI engineer inside who you basically say, okay, run this command and now the system is going to effectively manage itself and even give back feedback on things it's learning on site, but then it's all in the control of the company.
- 00:42:05 It's now their asset that's accumulating. We don't get any knowledge that's accumulated. We won't get any kind of ... Even telemetry they don't have to give us. And now



they're willing to expose it to everything because it's something that's secure in their barrier. So that's like the third kind of, I guess, third kind of leg of the stool of being able to have something where it's purely trusted because it's running inside, but we can do that now because it's an AI powered engineer inside that is managing the system as well as accomplishing the task at hand.

- Jon Krohn: 00:42:40 One of the key elements that I think allows that to work for you, allows it to work for Genesys as well as your clients, is something that you mentioned already earlier, but I want to highlight how important this is. You have a blog post that we'll link to in the show notes called How Genesis Automates Data Pipeline Development in Hours. And in that, you talk about how the agents escalate when confidence is low rather than force and answer. Tell us more about that and how important that is. Now
- Matt Glickman: 00:43:07 This was a big thing that we realized that we had crossed and one of these kind of massive steps forward was up until ... And this really happened when the reasoning models landed early last year, which everyone is excited about and basically show that you could scale up inference time thinking. But what came out of that was the ability for the models to be guided to be much more self-reflective on an answer. So you give an answer and up until then, you could coax them, you could threaten them with violence, which I still think is a terrible idea that'll be come back to haunt us when terminators come. But there's nothing you could do to basically try to really get them to say, how confident are you?
- Jon Krohn: 00:43:55 Yeah. There was an interesting study that showed that saying really aggressive things actually gets you like 10% more accurate response.
- Matt Glickman: 00:44:02 To me, it's just not worth the 10% to be on that list. But we saw that the real big win that was not often talked



about was this confidence indicator where you could basically like, how confident are you that you did everything I told you? And it would be very good at telling you, I'm 90% confident and this is why this is the one thing that I'm not confident about that I need clarification on. So that was a big kind of moment for us. And we basically now harness that. So every step of the way when we're going on these complex projects, these complex missions, we're constantly asking, okay, you did that, and how confident are you that you did it correctly? And again, if you are, great, show me the artifact. If you're not, then go back and try again. But when you're not, escalate what are the missing pieces, right?

00:44:54 And with those, that becomes like, well, I had to guess what this formula was because I just couldn't find it, which if you don't ask for, it's just going to be something that goes under the radar. Again, coming back to the human, we hold the humans to a higher bar, humans do this all the time. You basically just make these jumps of logic. You're like, "I think that looks right. I'm going to go with that. " And now someone asks you, "How'd you come up with that? " "Well, I actually made that up. It sounded good. "They do the same thing, but if you call them on it, you actually get it to be much more productive. And then they basically can ask you the intelligent questions. And then most important thing, and this is no thing more frustrating than when you're asked a question and you basically give the answer, you really want that to be applied next time.

00:45:47 Nothing's more frustrating without you dealing with a human or an AI where it's like, you asked me a question, I gave you the answer, you said that was a great observation. Don't come back to me tomorrow and say like, " Hey, what do you think about this? "It's like, " I told you that yesterday. "If you do that right and you capture those moments and those nuggets and you do it in a secure way, now there's no limit on what you can take on



and it compounds because now once you understand all this logic and how these businesses operate, you can then move up the stack because now you understand all the semantic, all the flows, all the kind of how we got here, and now we're finding that our customers are pulling us to go further up stack because, well, that's great now. Can you help me actually present that to the board?

00:46:37 Sure. And now we have our systems actually able to produce a well thought out presentation because it's grounded in the facts of how that actually operates.

Jon Krohn: 00:46:46 Makes a lot of sense. And to dig into that just a little bit more on this correctness point, you've previously in an interview stated how with consumer facing products, novelty is often one of the most important characteristics of AI systems. You're talking there about some kind of inventiveness that these models tend to have filling in the blanks, that novelty piece is key for consumer facing products, but for enterprise products, it's correctness.

Matt Glickman: 00:47:12 Yes. Correctness is everything grounded in truth and the ability to navigate these complex organizations and extract out what is correct. The best run organizations do not have the rule book that's like, this is all the things that are correct. They just don't. It's in people's heads. Some of it's encoded in the systems still to this day. Why do people stay at these regulated industries for such long careers because of all this knowledge that gets stored up here? It's cheaper to keep that person and keep on paying them than trying to download it out of them and put it somewhere else. Again, I was at Goldman for 25 years and that was normal. Still to this day, I mean, people stay there, this is going to change because all that knowledge is going to become what was a liability for companies is now going to become an asset where now if I can have ... Imagine, if I have this system that has all the knowledge about all of how a major banker or healthcare company, thousands and thousands of employees now have put all



this knowledge implicitly in a knowledge base that the company owns, I mean, you're going to see companies, you're going to see M&A that a variable of the M&A is going to be like, well, do they have a consolidated knowledge based context graph of how the firm operates?

00:48:41 I'm going to value that more versus trying to ... Because people are going to leave,

Jon Krohn: 00:48:45 Right?

Matt Glickman: 00:48:45 What happens every time we have M&A and it's like risk and people going to leave, who's going to keep and who we going to pay. Yeah, we got it all.

Jon Krohn: 00:48:53 It's like the classic of how you get a much better multiple on a SaaS product business relative to a consulting firm. Exactly. Because the consulting firm has so much- Not scalable. Human knowledge that leaves with the people. Yep. Speaking of these knowledge bases, you call them living context graphs. And so these are systems designed so that this institutional knowledge compounds over time and has never lost. And it looks like this is born from your experience watching critical organization knowledge disappear across data teams at Goldman Sachs and at Snowflake. So can you tell us more? Obviously you can't get into too much about your secret sauce, but we have a technical audience. I'm sure they'd love to hear a bit more about how these living context graphs work.

Matt Glickman: 00:49:36 No, and it's interesting that we fell into it. We didn't sort of say, "We're going to build this agent platform. We need to build a context graph." It was basically, we realized that the missing piece that we were constantly feeding all these agents to get going was a bunch of context. And that was the big human element was like, okay, human, go and find all the documents, find all the repos, point me all the relevant databases, and then it would go and be super successful. And we also, that's not going to scale



because no one would want us to want to do that gathering because now it's a big, you've now taken off some of the work, but you're making the humans do more other kind of hard work. So what we basically do as part of Genesys onboarding is you connect this system to all your databases, all your repos, all your SaaS tools or as many as you want to, or on- prem, in cloud, wherever, and then they go about and start crawling.

00:50:37 So think about a traditional kind of web crawler, but in the context of wanting to understand all the data relationships that exist amongst organization, a spreadsheet here, a database there, an API call here, like all the things and then effectively layer on top of each other. So now you can see that this code is referencing these tables and these APIs, and now you can effectively build up a graph and you can see it, we have a demo on our website, you can see it how it literally becomes almost looking like a social graph, but of the data relationships amongst the firm. And it gets super complex super fast. So the goal is not to have any human ever get their head around it, but an AI, again, loves this stuff. And the crazy thing about it was that we did. We did this crawl, we got all this kind of built up with this graph with all this kind of metadata on it, and we just gave our agents tools to navigate this graph without even explaining why.

00:51:36 We've since now kind of guarded it, I mean, guardrail a little bit, but it instantly said, "This is great. I can now understand." And it was like we'd given it the secret formula and it basically just started crawling and say, oh, for example, a user wants to add a new column to a table on a report. And our agents basically went and Eve is our master agent, Genesis Eve, Adam,

00:52:07 Them clever at the time, but basically Eve will go and say, "Okay, well, where am I going to get this data from? Let me see where there's similar data around." And it goes and figures like, "How did I produce this report?" So I



understand that kind of path. It's like, can I follow that same path back to its source and was there a missing field? Yes or no? Maybe. Or is there some other field or some other reports similar that I can find and basically search on this multidimensional graph space and try to find a similar, semantically similar search on the space and find something similar that I can connect to and then pull that data through and build the pipeline. So it just fell out of that and the AI was, it was almost like it was meant to be, but this was going to become, and now it's been self-fulfilling where now everyone's figured this out and it's getting better and better knowing how to navigate this.

00:53:06 So yeah, and that basically it was a huge unlock for us and then makes our systems come online that much faster, that much more productively and it removes the overhead for humans to kind of train the system or effectively similar to how Google went and crawled the internet. They didn't have to ask people to go to Yahoo and put in links and maintain that. They just did it passively and uncovered all that relationships on the fly doing the same thing, but we're not doing it to produce a graph. We're doing it to do better data projects and data engineering. And it's because of that reasoning, people are willing to kind of connect more and more systems because they see the value.

Jon Krohn: 00:53:53 Yeah. Yeah. You create the graph to enable a better AI system to be able to crawl that graph and have knowledge more quickly, better, more concretely represented ideas. Yeah. Cool. All right. So if people are listening and they're thinking, "I'd love to have this personally- It's obvious.

00:54:12 Or in my organization," a key part of the adoption problem here is that it seems like a lot of people are still thinking about the pre-February 2026 mindset of, is this an AI use case instead of why shouldn't this be done with AI? And so you outlined a four-phase adoption model that



begins with assessment and ends with scaled autonomy. Do you know what I'm talking about?

Matt Glickman: 00:54:41 Yes.

Jon Krohn: 00:54:41 Can you tell us about that? Yes,

Matt Glickman: 00:54:41 Yes. No, I think you have to be thinking about how you're going to get there and what are the steps along the way. You're not going to just turn the system on and it's going to be fully autonomous, nor would you want it to be because you won't ... It's actually more for us humans to come along on the journey than the AIs who may be ready for it before we are.

00:55:11 So it's understanding the problem, being able to kind of understand what would a human do to solve that problem, and then be able to get there in a way where when the AI is successful, we understand how it's achieved it, and then we're willing to let go. And I think that kind of ... But it's really more about the humans being comfortable about it, but you have to constantly be trying to push the limits because this space is evolving faster than we all thought, but also don't be fooled by these amazing coding agents who are, again, who inspire us to do things, but without the right guardrails and without the right kind of contextual understanding, they can cause more damage than good. But with the right guardrails and the right context and the right human oversight, it's a wonderful time to be in this space.

00:56:24 Yeah.

Jon Krohn: 00:56:24 How do you convince your enterprise clients, these big organizations with lots of liability risk and lots of people internally who are probably skeptical of what AI agents are capable of doing today, how do you convince them to make the leap from chatting continuously with a conversational agent to delegating to a team of agents



where say Eve, as you've described it, is your head agent kind of orchestrating, how do you convince your clients that that kind of delegation and trust is the right time?

Matt Glickman: 00:57:04 By showing, not telling. At the end of the day, the pain is so high and these places demand is so high and been so underserved for so long that the answer is always like, if this works, it's a no-brainer. And doing it in a way that it's doing things, it's all audited, it's all documented, it's not like running rampant. It's going to follow your normal processes, right? It's going to test things. It's going to run in development. It's going to provide a code review and a PR to submit to your CICD pipeline. It's no worse, and I argue better than you hire a new employee because you're going to have processes that prevent that new employee from going and crashing production, right? It's basically as much as risk as that. So the only thing they have to lose is that it might not work, but they all have such long backlogs.

00:58:07 They have so much pressure to do more with less. And almost everyone sees this opportunity to basically now get out from behind the curtain and say like, let me show you what this can do. I can actually focus on a business impacting goal instead of working on this machinery. So this again, the pent-up demand, but the risk is as they already have processes in place that prevent rogue developers from doing wrong. Eve just signs on as just another developer on the team. So the only obstacle we've had is that people just don't believe that it works because either they've had their own experiences or they've tried one of these coding agents or they've worse tried to build their own agents, which I think is also a fallacy, like focus on problems, focus on outcomes. Traditionally, you didn't hire a consultant and figure out, can they give them tests on how they can ... You basically like ... You said, "I need you to do this migration." And they said, "Okay, well, this is what it's going to cost, and this is the people, and this is how we're going to do it."



00:59:24 " And you would compare different options, and you didn't care how they were going to do it, you didn't care which people they're going to use, and you were selling outcomes. Similarly, that's what businesses want, and enterprises are getting even more now kind of critical where if it's not core to their business, back to my earlier point, they know doing it for revenue will win, so why not pick a winner instead of trying to keep up with something that is just going to be accelerating out of their reach?

Jon Krohn: 00:59:58 Yeah. Got to stay on top of this for Fast moving thing, for sure. You described two moments in your career, first with Snowflake and later with Genesis. When you realized you could either sit around and hope or you could actively help shape what came next, how do you know when a technology shift, like the one that I think we are both convinced we're in and hopefully a lot of listeners as well, how do you know when that technology shift is real enough that you should stop analyzing it and start building for it?

Matt Glickman: 01:00:28 Yeah. No, I think I truly believe we're all here on this planet or wherever planet we end up going to for a purpose. And if you're kind of self-aware enough, you kind of know what you're here to solve.

Jon Krohn: 01:00:46 To unleash the machines.

Matt Glickman: 01:00:47 Unleash the machines or just do it, but more even a meta problem. For me, my entire career has been about unlocking this limitation of there's not enough people who can actually understand technology, understand a business problem and kind of connect it to. And if I look at a higher level, I've been trying to solve the problem forever. Build with platforms, with going to Snowflake and trying to provide that as a capability other people could use. And now at Genesis basically unleashing that with AI, I think just taking a step back in your day-to-day to understand what is the world trying to tell me? And



you're going to have these moments like my meeting with Benoit when they came to Goldman, or when I had my moment where the early GPT-4 basically explained to me before anyone else was talking about agents, that it could call functions on my behalf, or I can call functions on its behalf as a way to basically what agents became.

01:01:56 Being aware of these moments and saying, again, the customers I was talking to who were constantly saying that they were trying to basically democratize their data teams. Be aware of what is the world trying to tell me here? And do I have an earlier view on where this is going to go than the rest of the world? If I'm behind, that's not the time to jump in and start doing this. If you're ahead and you know this is in your wheelhouse because of something that you've seen that maybe the rest of the world hasn't seen, that's likely the time to jump in. The challenge is that these moments are now kind of, the openings are smaller now because of the pace that we're in. We are in the event horizon. Anyone who doubts it is clearly not touching the space. It is exactly like Kurzweil predicted, we're in the exponential, it's accelerating.

01:02:56 You can no longer track the improvements. There's no plateau happening anytime soon. Scaling is continuing to scale. So the only challenge in this kind of approach is that you have to be more aware. You used to have openings and you'd see them and they would present themselves and you have time to think about it. Now it's going to be things where it's going to happen and it's going to be like, instead of being like a month opportunity, it's going to be a much shorter timeframe. But I think if, given what you can now do and build and try in the shorter periods of time, I mean, you can in a weekend, right? The guy who built OpenClaw did it in a weekend, and that changed the entire game on personal agents. So I think trust your instincts, but always ask yourself, why am I seeing this before other people?



- 01:03:50 Because if you can answer that question, then you should jump in.
- Jon Krohn: 01:03:54 Really cool. Great guidance there. Thanks so much, Matt Glickman for that guidance at the end and some excitement, anxiety. How can we get on top of this so quickly? Yeah. Our brain registers the same neural response for excitement and anxiety. And then it's up to your cortex to interpret that sensation.
- Matt Glickman: 01:04:22 It's pretty cool.
- Jon Krohn: 01:04:23 And hopefully most of us are taking a step back and taking the opportunity that we're in this event horizon and there's lots of exciting things we can be doing. Let's take a step back, use this incredible tooling to build an open claw type thing in a weekend because you can do it now. You can. And make a huge impact. So yeah, whether it's with adopting Genesis as an AI data engineer within your organization or building something yourself, very exciting times. Thank you so much for sharing so much knowledge that you've accumulated over these decades. Sure. Really appreciate it, Matt. Before I let you go, I always ask my guests for a book recommendation.
- Matt Glickman: 01:05:02 No question. Hitchhiker's Guide to the Galaxy.
- Jon Krohn: 01:05:05 Nice.
- Matt Glickman: 01:05:05 By Douglas Adams. If you haven't read it, you should read it immediately. If you've read it, you should read it again. It is uncanny how the entire AI explosion we're going through was predicted with such precision in a comical way. And it's basically, they try to build a supercomputer that's the super AI. And of course, I won't spoil the punchline, but yeah, it is exactly what we're living through what Douglas Adams predicted.
- Jon Krohn: 01:05:40 Yeah. I think I can make this point without giving anything away, but something that's very different about



the supercomputer that they're building there is that it takes a very long time to compute, to do the big inference. Whereas it seems like something very different about what we're going through now.

- Matt Glickman: 01:06:01 Don't assume we're at the end of the book yet. Because think about what we're trying to do. If you even listen to what XAI is trying to do with Groc. I mean, there it's like imitating reality because they're actually thinking about that book a lot. Elon is when he's funding XAI, but what we're doing now is it could be the early stages of the big computer that is built. We have not gotten to the point where any of these systems are actually discovering new things yet. So I think we're really early in that buildup. Those you didn't catch it, I think it was last weekend, Kanuth, who's one of the famous computer scientists who basically kind of disappeared after basically defining how we should do computer science. He's like 90 years old and he just published a paper that he co-wrote with an AI about a mathematical problem that he had not seen a solution for yet.
- 01:07:07 So that's like maybe one of the first examples that we're approaching this new place that the book
- Jon Krohn: 01:07:14 Plans
- Matt Glickman: 01:07:15 About.
- Jon Krohn: 01:07:15 I think OpenAI researchers have been talking about 2026 or 2027, making kind of new physics discoveries they anticipate.
- Matt Glickman: 01:07:23 It's going to happen. And one more interesting, I'll leave the audiences with some most interesting experiment that I've heard that is going to be done, which is basically to roll back all the training data to what was available to Einstein at the time, but not anything else that was published in science or anything else. And then see if that



kind of time traveled model can produce the theory of relativity. That's going to be the ultimate test of- It's a fun idea. Yeah. It's super interesting. And I mean, I don't know. It seemed pretty hard to do to isolate, like no written word, no newspapers or nothing, but if you can do it, I think that's going to prove that we're

- Jon Krohn: 01:08:07 There. Sure. Yeah. Leakage could be a key problem. Leakers is a
- Matt Glickman: 01:08:09 Problem. Yeah,
- Jon Krohn: 01:08:10 Exactly.
- Matt Glickman: 01:08:11 Interesting times though. So awesome.
- Jon Krohn: 01:08:13 Interesting times for sure. Matt, for people who want more of your insights or more information on Genesis after this episode, how do they follow you?
- Matt Glickman: 01:08:20 Yeah. So website, [genesiscomputing.ai](http://genesiscomputing.ai), also.com, which was an interesting purchase. And I'm on Twitter at Matthew Glickman and on LinkedIn, Genesis Computing or Mac Lickman. There is a Doppelganger out there where I finally actually crossed paths with, I am not the West Coast Mac Lickman. I am the East Coast Mac Licket.
- Jon Krohn: 01:08:42 I think if I remember correctly on LinkedIn, you're Matthew J.
- Matt Glickman: 01:08:45 Yes.
- Jon Krohn: 01:08:45 I am Matthew
- Matt Glickman: 01:08:45 J.
- Jon Krohn: 01:08:47 Just
- Matt Glickman: 01:08:47 To have some separation.



- Jon Krohn: 01:08:49 Yeah, there were a couple times before we booked you for the episode where Natalie on my team showed me that other, the double ganger one. Yeah, because he's in tech as well, right? And I'm like, "Is that him?" No, that's not fun.
- Matt Glickman: 01:09:01 Yep. I'm the other guy.
- Jon Krohn: 01:09:03 Nice. All right, Matt, thank you so much for coming to record with me in person. This was a really interesting episode. Yeah, thanks for having me. Really exciting
- Matt Glickman: 01:09:08 Times. Awesome. Thanks a lot.
- Jon Krohn: 01:09:13 Lots of food for thought. In today's episode with Matt Glickman in it, he covered how February 2026 marked the moment the latest frontier models crossed a threshold where they could handle complex multi-step data engineering workflows that previously required human expertise. And this big change means there's no going back. He talked about also how finance and healthcare were late to adopt the cloud, but are among the earliest and most aggressive adopters of AI. How Genesis computing deploys its agentic platform directly inside a client's environment, more like onboarding a new employee than adopting a SaaS product so that all accumulated knowledge remains the company's asset. And he talked about how rather than acting as a copilot that waits for human instruction step by step, Genesis inverts the model. Agents work autonomously on complex data engineering tasks, only escalating to humans when their confidence is low, memorializing every answer so they never ask the same question twice.
- 01:10:08 As always, you can get all the show notes, including the transcript for this episode, the video, recording, any materials mentioned on the show, the URLs for Matt's social media profiles, as well as my own at [superdatascience.com/981](http://superdatascience.com/981). All right, that's it. Thanks to



everyone on the Super Data Science Podcast team, our podcast manager, Sonja Brajovic, media editor, Mario Pombo, our partnerships team Natalie Ziajski, our researcher, Serg Masís writer, Dr. Zara Karschay, and our founder Kirill Eremenko. Thanks to all of them for producing another stellar episode for us today for enabling that super team to create this free podcast for you. We're deeply grateful to our sponsors. You can support the show by checking out our sponsor's links, or if you'd ever like to sponsor an episode yourself, you can get the details on how by making your way to [jonkrohn.com/podcast](http://jonkrohn.com/podcast). Otherwise, please help us out by sharing this episode with people who would love to hear it.

01:10:59 Review it on your favorite podcasting app or on YouTube. If you write a written review on Apple Podcasts, I will read that on air in an upcoming episode. Obviously subscribe if you're not already subscriber, but most importantly, I just hope you'll keep on tuning in. I'm so grateful to have you listening, and I hope I can continue to make episodes you'd love for years and years to come. Till next time, keep on rocking it out there, and I'm looking forward to enjoying another round of the Super Data Science Podcast with you very soon.