Talks at GS

How Aaru Uses Advanced AI Simulations to Predict Human Behavior

Date of recording: July 22, 2026

Cam Fink: We help you understand the composition of any population on the globe, and what decisions they make. What marketing campaigns they're going to click on, what subway lines they're going to ride.

Ned Koh: Our model is able to process how humans are able to think.

[MUSIC INTRO]

Marco Argenti: Welcome to Talks at GS. And also welcome to Ned and Cam from Aaru. Nowadays with AI almost like anything is possible, including being 19 and then, you know, having a $1 billion valuation startup, which is pretty incredible. So, since then, Aaru has become one of the most closely watched companies in the research and intelligence space. Thank you for being here.

Cam Fink: Thank you, Marco, for hosting us. We're

really glad to have the opportunity today to talk a bit more about what we do.

Marco Argenti: The first question I'm going to ask you is, when was the first time you kind of realized that AI was something big?

Cam Fink: So, it's interesting you ask this because the ChatGPT moment was right in our senior year of high school. So, I remember, you know, having worked in the NLP space, in the computer vision space, where Ned and I were researchers before ChatGPT was released. You could tell that like AI was gaining a lot of capabilities very quickly.

But it wasn't until it became actually productized into something consumers could use that all of a sudden it felt like not only was AI going to gain capabilities, but actually that it would create serious change in how people work on, you know, ideas, on products. But even still, like I remember the most enterprising students were already using AI to evaluate their answers to papers or come up with new ideas or test new things.

And if you were even more enterprising before that, you would have been using the GPT-3 API to try and finish papers. So, you know, I think that was truly the moment that it became clear AI was going to be inevitable.

Marco Argenti: Was there another ah-ha moment recently? Because, you know, right now it feels like things are changing so fast. What's another inflection point that you've seen?

Ned Koh: The tools are incredibly powerful for being used for incredibly rudimentary application: note taking, you know, or summarization as like the key use cases or demonstrations of like how this can help you. And I still think in many businesses it is. It's easy to get behind. There's not really any friction to adopting those technologies. And so, I think that the larger change is this wave of product-first AI businesses. You saw that in health care. You saw that in, you know, financial services and beyond. But I think it was when it went beyond just those into more complex applications, for sure.

Marco Argenti: So, rolling back, like, kind of two basics. If you're in front of a CEO that doesn't know what Aaru is

and you have a few seconds to explain, what would you say to them? What's your pitch? Like you're 30 second pitch.

Cam Fink: We predict human behavior. So, we help you understand the composition of any population on the globe and what decisions they make. We can help you understand how U.S. household purchase decision makers are going to purchase products, what marketing campaigns they're going to click on, what music they're going to stream, what subway lines they're going to ride. That's really the core of what we do: predicting behavior to make better business decisions.

Marco Argenti: And I think we met maybe two years ago for the first time. So, really early on. So, what have you learned in terms of sort of what kinds of decisions is AI really good at predicting and what others are not that good?

Ned Koh: Yeah, I think the first place that you look when you define what good means is understanding like the relative for how people make those decisions today. Right? So, let's take something that's like very simple - product launches. So, you have someone who's going in

front of a survey, you're asking people questions. It's easy to benchmark. And that's a very interesting market. Product innovation is massive for us.

But when you think about like use cases and how that's kind of diverged across the board, the most relevant ones for us and where we offer the most amount of value, of course, are places where you may not traditionally have data. Right? And so, that may exist in a population who doesn't respond to surveys. The example I like to give is like everybody in this room is wearing a shirt. Did anyone tell someone how much they were willing to pay for it? Probably not. Right?

But then you get into like other categories. Right. You know this answer better than I do. But you know, there are four industries that make up about like 15% of the GDP of the U.S. Airlines, telecom, utilities, insurance. What are the times that matter to those companies about their behavior? When do you remember who your cell provider is? When your cell service goes out. When do you remember that you pay bills to Con Ed? When your power goes out. When do you remember that you fly a specific airline versus another? When your flight gets delayed. Right? And so, on

and so forth. When do you remember who your insurance is? When you crash a car.

And so, those businesses, despite being massive, the most important place and time for them to understand the behavior of their customers is somewhere where it's fundamentally impossible to get inside their head inside that moment. Right? And so, the most interesting use cases to us have always been those areas that have the most challenging to reach groups. The time frames of traditional research, it can take weeks and weeks or months to get people in a room and ask them questions or show them a product. You don't have the ability to do that. Right?

So, things that are compressed in timeline are the use cases that are most relevant for simulation because human insight just doesn't reach that same level.

Marco Argenti: But why can't you just ask ChatGPT?

Cam Fink: Language models themselves alone can't understand human behavior with the depth and the capability that is required to actually allow a business to

trust their decisions. There are a lot of startups out there now doing what you might call synthetic personas. That's sort of like the last generation of the technology we build. They can be useful. Those startups can help you get like a really quick read on marketing copy. But if you want a decision that you can trust, you need something that's grounded in real world outcomes. And that's where someone like Aaru would come in.

Everything we do is AI. Right? So, what underpins our capabilities to predict human behavior are models. Models that we train. Models that leverage LLMs but aren't LLMs alone. But through fusing kind of LLMs with statistical models and traditional mathematical models, we're able to be far more accurate and take advantage of all of LLMs' advantages without falling into those kind of gaps with LLMs and predicting behavior.

Marco Argenti: And how do you verify the accuracy and the results? I mean, what is what's ground truth for you guys?

Ned Koh: At the end of the day, if you're trying to predict and make a decision better, what you should be

aiming for is a better business outcome. That was very first principle to us, right? You want to benchmark it to the sales result, to the product selling better or selling worse or the limited time offer not flying off the shelves like you wanted it to.

The way that we've been able to grow our business so quickly is by making sure that people walk out of a pilot with us and don't just say, "Yes, we came to a decision faster," but rather we came to a decision better than we otherwise would have.

Marco Argenti: So, let's actually kind of bring it to life a little bit with some examples. What does the client actually get? Give us some examples of that.

Cam Fink: So, we do a lot of work with CPG companies. We do a lot of work with retail, beauty, apparel, kind of like anything that touches a consumer brand. And every kind of business ultimately sells a product. But for consumer brands, it's typically really linear. You can hold the product in your hand, right? A canned beverage, a carton of milk, whatever it might be.

So, we do a lot of product innovation work where we'll go work with, you know, a canned sparkling water company, for example. And that sparkling water company will come to us and say, "Hey, we've built this great business around this one core product. What is our next new product line, right? Is it still hydration? Is it a shelf-stable dairy product? Is it an energy drink? Is it--," you know, all of these things.

And we can use Aaru to do everything from the very core conceptual stage of what is the next frontier, right, all the way into figuring out who exactly is the purchasing audience. Then into the product concept itself specifically. So, say we pick out that it's an energy drink. We figure out who the energy drink purchaser is going to be from this brand. And then we figure out, okay, what flavors of energy drink do we want to put out in the world? And then we can do everything from getting as niche as optimizing the caffeine content to helping figure out the right pack size and carry that entire product innovation cycle from start to finish in just a few weeks.

And one of our clients we worked with, their product innovation cycle traditionally was nine months long, right?

That's actually pretty quick for the CPG industry. A lot of the big global companies might be taking 12 months, 18 months, to go from product idea or, you know, new market entrant to ready for manufacturing. We managed to squeeze it down into eight weeks. That was a really, really, really powerful case study. And now the product ended up hitting the shelves earlier this year. And they saw more success than they had with any of their previous product launches before.

Marco Argenti: And if I remember, that you guys can decide the size of the core. Tell us a little bit about that spectrum.

Cam Fink: Yeah, we don't go down to individual-by-individual level, right? We're never processing PII or any kind of individual data. So, we can't recreate Tom Smith from Iowa or Cameron Fink from New York City, right? We can get down to audiences that in the real world might only be a couple hundred people.

I remember at one point doing hospitals that have to purchase this third-party risk management software. You're probably talking only a few hundred to low

thousands of people nationwide. And yet we were still able to recreate that audience. And for the company that we did do that study for, they saw a massive impact from the work we could do together.

So, we're getting really niche. We can talk about very small groups of buyers or social media influencers or Gen Z in one particular city metro. And then we can go as wide as US household purchase decision-makers or British residents or, you know, even residents of the EU if you want it.

Marco Argenti: And where do you get your data?

Cam Fink: I mean, it's a massive wealth of data sources, right? So, when we talk about how we leverage that data, we have audience data and we have behavioral outcomes data. On the audience data component, that's what helps us understand who actually exists inside of a population. So, when you come to me and you say, "Hey, I want to talk to ultra high net worth individuals in, you know, the United States and Canada," we can actually tell you what the distribution of that population looks like.

And then we have behavioral outcome data. That helps us figure out what decisions are they going to make. So, that would, you know, be for each of these given value propositions, how will UHMW individuals react to them?

Now, when it comes to where that data actually comes from, you can think about three different buckets for each of the data sources that we leverage. The first would be publicly accessible data. We're leveraging tons of things like the census, academic research where there's license to train on top of it, you know, anything else available publicly that our model is able to use.

Then there's a lot of data we purchase. So, we buy data from everyone from credit card, you know, purchase history data providers to receipts, there are music streams, social media data. And then I would say really the final component is proprietary partnerships that we leverage. So, we do a lot of work with organizations that wouldn't fit the shape of traditional customers of ours but still want access to the impact of our software. We have done data sharing through those sorts of engagements before that's been incredibly powerful for what we do too.

Marco Argenti: How do you make your clients comfortable that they can share their most valuable data?

Cam Fink: Every client is sandboxed entirely. We're SOC 2 compliant. I didn't expect that I'd be saying that up on stage. But just in case anyone was worried, we have any InfoSec people in the room, we are SOC 2 compliant. We have all the compliance certifications we can. Most importantly, we do not train on customer data. And all of our clients are sandboxed.

So, we actually end up being trusted as one of the most secure places to store, you know, that data because you would never be able to leverage a traditional research provider for anything sensitive. Right? You wouldn't want to go survey people and then someone takes a screenshot and says, "Hey, you know, my favorite brand was polling me about this on the internet." Instead, when you're leveraging our software, it's kept entirely secure to, you know, the model that we train.

Marco Argenti: You just said the models that we train. So, has your attitude and approach shifted over the years with regards to kind of using your own models versus

using frontier models versus using quasi frontier open-source models?

Cam Fink: Yeah, our approach has definitely changed. I mean, when we started the business, we were definitely more in this realm of like, hey, can you take an LLM and train it on top of survey results? We found some of our own issues with that, right? Historically, we found it wasn't extensible. And that's what really led us to the next generation of our technology.

So, those are all the models that we built ourselves. And that was really the big leap forward in terms of how our technology was able to reach the accuracy level that it is today.

And so, we've been on, you know, all of our own models ever since that point in time. LLMs help us write analysis. They help us do tons of tasks as part of the simulation pipeline. But LLMs aren't driving the fundamental underlying prediction and behavior themselves.

Now, when it comes to the LLMs we integrate, we use tons of different providers. We default to whatever is the best

provider for each task that we have to do.

Marco Argenti: Which could be, you know, somehow maybe filtering all the data or explaining or prompt interpretation. But then the actual prediction machine or the simulation machine is your own. And that's your secret sauce.

Ned Koh: That is the secret sauce. Every single economist in the history of like behavioral science and everything else that you can go into has dreamed to be able to build a model off of something that's based on outcomes. You want to be able to remove as much bias as possible. Right? And so, not always is the magic and the data exclusively. It's much more in how our model is able to process things that are in many ways useless by their own right into something that can become a holistic view of how humans are able to think.

Marco Argenti: If every customer is protected and everything, how do you evolve the model?

Cam Fink: We're integrating new data sources all the time, right? So, the data that we have alone is enough that

we are able to go out and make tons of valuable predictions. We can generate a lot of insight without having to integrate client data. Integrating client data for us is the key to unlock, you know, much deeper capabilities and integrations, both in terms of, you know, the simulation component and also in terms of how simulation results end up getting used.

Marco Argenti: I get that. But within, let's say, client engagement though, let's say we use it in private wealth in GS and then we want to kind of continue to learn, so do you allow your clients to continuously learn on the model that you provide and that stays theirs?

Cam Fink: It's actually interesting. A huge part of what we're talking about now is this idea of compounding intelligence. One of the biggest issues that our clients feel, one of their biggest pains, is this idea that traditional insight is very static, right? You run one survey, you run another survey, maybe you run a quarterly brand tracker, but it's not like those survey results are propelling forth the next survey that you run.

And when you think about simulation, what's really

incredible about the difference with our platform is that it's not just the power to like access more accurate, faster results. Our simulations can also contain the insights from all the simulations before it. And what that means is that it actually builds and compounds in value over time.

You can imagine someone using our platform at the end of running nine projects about studying the Gen Z audience together, could come to us and say, "What are the insights, you know, we didn't even know we had about the Gen Z audience?" And that can go into helping build the next product that people might be interested in. So, a big focus for us is that compounding intelligence piece.

Marco Argenti: Are you guys mostly focused on the US, or do you have a lot of international customers?

Cam Fink: Yeah, we're a totally global business in terms of what we support. So, when you think about the range of Aaru today, it's over 130 countries across the globe. All of the EU, all of North America, much of South Asia, Southeast Asia, you know, and beyond. In the future, we are actually planning international expansion starting soon. But today, the bulk of our customers are based in

North America.

Marco Argenti: And so, on that note, so do you think you need to tune sort of your models by culture and by countries?

Ned Koh: I think one of the best markets that is super relevant for a lot of the biggest enterprises that we work with today is expansion into Germany is really interesting, right? I've seen this across customers of ours in consumer electronics. I've seen this across customers of ours in financial services, right? Like it really goes across so many areas.

The amount of safety regulations that people care about in like South Florida versus in, you know, Central Germany are vastly different. I can tell you that, right? Like the reality is, you know, you talk about why do we train models within this time when you're actually in runtime? Why do we keep them updated? These sorts of elements. The news, the way they digest social media, the things that changed their environment is what we like to call it within a given sim. Absolutely do. The context of how you and I make decisions on today versus tomorrow, absolutely do change.

Does my race change? No, that's not the most important part to measure. So, we absolutely have to do that and have to do it on a hyper localized basis.

Marco Argenti: What's your biggest bottleneck at this point?

Cam Fink: 100% hiring. Like it's the best people in the world is the number one thing you need for any fast-growing business. And so, when we think about what is going to-- has defined the last six months, going to define the next six months, next six years of Aaru, it's all going to be about working with the best people in the world.

Marco Argenti: And what kind of people are you looking for?

Cam Fink: The first thing that we evaluate for is raw capability. We do look for people who are the best in the world at something. Right? People who are truly excellent at the craft or at their craft, so to speak, and represent the pinnacle of understanding knowledge capability in a given area. That's something that every person wants. Right? I'm

sure Goldman wants people who are the best in their craft as well.

We also look for people who are really high slope. So, even if you're not the absolute best in the world at something today, do you have the potential to be that in the future? Can you learn quickly? Can you adapt is something that's really important to us as well.

And then finally, we look for people who are real collaborators. Every time someone joins our team, it's not just about whether or not they themselves are a really, really brilliant individual. Tons of people who are super smart and, you know, would be lucky to work at any incredible business on the globe. We spend a lot of time focusing on who is exactly right for Aaru and the people who have much to teach our team, much to teach myself, but we also have much to teach them and they can be, you know, leverage, exponential or multiplicative leverage to the people around them.

Ned Koh: Most of our research team, right, you'll have people who are pure math backgrounds. You'll have people who have done Alzheimer's and dementia research, right,

who are able to tie in these really interesting areas. On the other side, we've recently brought on the head of research and data science from Qualtrics, who was there for over 13 years. Right? It's like you need some of those incredibly experienced perspectives to understand how to solve these problems. But you also need very fresh ones to think outside of the box to be able to do it as well. It's all about a combination and a balance.

Cam Fink: I think what we believed was true when we started this business is still true for the most part. Right? Like, we've always believed in being open communicators. We've always believed internally in, you know, the power of learning from each other, the power of collaboration. I think those are all important things. We've always believed in wanting people who are excellent at their craft.

When we started to think about it, it's that as a growing business, it's actually time to write that down for the first time and distill it into the things that matter most for us. And that is what led us to kind of starting to think more about culture and principles in the last few weeks than we had previously.

Ned Koh: To that point, a really interesting thing that we've been discussing, especially with some of the larger implications of being able to model behavior, is this idea of like trustworthiness of a model. There's a very big difference between being worthy of trust and being a trustworthy company. And we certainly try to stay on the prior, right? We define that as being worthy of trust to the people who work next to you, right? That you're going to do an incredible job. It's being worthy of trust of your investors. It's being worthy of trust of your customers. It's being worthy of trust of the world that you're using the technology for the biggest impacts you can.

Marco Argenti: Ned, Cam, thank you so much.

Ned Koh: Thank you so much for having us, Marco.

Marco Argenti: Was great. And good luck. Thank you so much.

Cam Fink: Thank you. Thank you.

The opinions and views expressed herein are as of the date of publication, subject to change without notice, and may 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, legal, or tax advice, a recommendation from any Goldman Sachs entity to take any particular action or be used as a basis for any other investment decision, or an offer or solicitation to purchase or sell any securities or 43 financial products. Any forward-looking statements, case studies, computations or examples set forth herein are for illustrative purposes only. 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 sponsorship, affiliation, endorsement, ownership or license rights between any such company and Goldman Sachs. 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.

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.

© 2026 Goldman Sachs. All rights reserved.