

Privately Speaking Podcast

May 2023



Shivani Sopory:

Hello, my name is Shivani Sopory. I am a partner at KPMG in our Private Enterprise group and I'm really excited to be guest hosting today's Privately Speaking podcast and to have with me, Elaheh Ahmadi, who is the CEO of Themis AI. Welcome, Elaheh.

Elaheh Ahmadi:

Thank you Shivani. As you said, my name is Elaheh Ahmadi. I'm the CEO and cofounder of Themis AI. A little bit about me, I was born and raised in Iran. I moved to the US after I got into MIT for my undergrad and I got my bachelor and a master degrees from MIT in computer science. And for my master's I was working with my now-cofounder on a core technology that we spun out and started Themis AI based out of.

Shivani Sopory:

Thank you for the quick intro. We're definitely going to get into a lot more about your background and Themis itself and really what your company does, which I think is a great way to move into the next topic. Can you tell a little bit about your origination story for your company? I know it's fairly specific because you did start research and you started out at MIT, so it'd be great to understand what your research was and how that turned into the company you're running right now.

Elaheh Ahmadi:

Yeah, of course. The research was done at Professor Daniela Rus's lab at MIT Computer Science and Artificial Intelligence Lab. And this is the work that we've done for the past six, seven years on what we used to, or in the lab, called it a debiasing technology. So basically what motivated us was one of the biggest challenges in machine learning application, especially in dataset, is that a lot of the times you have rare events or underrepresentation in certain groups or certain scenarios, but you want to train a model that still works well on those rare events because you want your system to be able to perform well under all circumstances, but if there isn't that much representation in a data set, then your model's not going to learn it. So we were very passionate about this and we were like, "Is it possible to solve it?"

The first area that we tried to tackle this problem was in autonomous vehicle. So one of the challenges there, again, is in the data set. Most of the data set is actually from a car driving down a highway. And that's good, when you use it for training, the model learns how to drive down a highway, but those are not all the scenarios that you drive. There is bad weather scenarios or the complex roads that you want the

vehicle to still be able to operate. But if they're not represented, how would the model learn it? So we came up with this algorithm that basically was able to automatically pick up on all of these rare events and make sure during training the model is seeing those more. So even though there isn't that much representation, the model is focusing the training time on those areas and that resulted in being able to outperform all other state-of-the-art models on autonomous vehicles.

That led into a next step. The team was, "Okay, we tried an autonomous vehicle," and that was around the time that there was a lot of issues with facial detection, the Google model, and Facebook model that they had to stop it or ban it. The team applied this technology on there and we were able to basically do the same underlying technology, be able to understand the dataset. So what we do is basically try to understand the underlying distribution and the underlying features that are in the dataset and use that to understand what are some areas that are underrepresented and make sure that the model learns those areas more.

One of the challenges of facial dataset is that there aren't that enough representation of racial minorities or people of different colors and that resulted in the model not being able to perform on there. Again, similarly, we were able to improve the model performance, especially on underrepresented groups and now perform all the other existing models. And we're like, "Okay, cool. We did it in two applications, let's see if it works on other data modalities," because until now it was only on images and that's when I joined the lab and I started working on adapting this technology on healthcare application.

In that scenario, we were working with tabular data and it was clinical trials data. The goal was to basically be able to predict whether the clinical trial was going to be successful in phase three given information on phase one and phase two. Similarly, we were able to again do the same thing, uncover some biases that were not obvious, especially in tabular data, it's a lot harder because you don't have anything to see, it's numbers and you need to make sense out of it. And again, we were able to outperform the other state-of-the-art models and this brings us back to summer of 2021 and the founding team got together and we were like, "Okay, we have this amazing technology and it seems like the industry needs it, so why not spin it out and help the industry with this challenge? So that was the very long version of how Themis was born.

Shivani Sopory:

Oh, I love that. It's very clear that you were almost testing a theory and then testing it a couple times over and it kept proving that your system was just better than anything that was out in the market, which is the best way to determine if something needs to actually enter the market.

Elaheh Ahmadi:

Yeah, exactly.

Shivani Sopory:

I'm just curious, how long were you all in research and development?

Elaheh Ahmadi:

I joined the lab almost end of 2019, but this was the research that was in the making since 2016, I believe.

Shivani Sopory:

That's pretty impressive. It's a long runway.

Elaheh Ahmadi:

Yeah. I think it shows the complexity of the problem that after that many years of research, we still have figured out maybe two parts of what actually makes a trustworthy AI model and there is a lot more research that needs to be done because these models are complex. There is a lot of things about them that we still don't know why they're even working to begin with. So being able to remove that challenge of risk out of it makes it a lot harder.

Shivani Sopory:

You mentioned trustworthy and I also noticed in your mission statement that your mission is to enable trustworthy AI. So for the audience, can you talk a little bit about what you see as what does trust mean to you and why it's important?

Elaheh Ahmadi:

There's a lot of different pillars that contribute to trust and making a model or a software trustworthy for a system. And for us as a team, we focus on three main pillars. One of them is fairness, so is your model going to perform well on all the subgroups that you have? And this can be social fairness or, going back to representation, representation fairness.

The second one is reliability and robustness. So you want your model to be reliable and be robust to new scenarios and different scenarios. And the last one is explainability. So you need your model to be able to tell you to some extent what's going on and tell you when it's failing and, when it's failing, what was the reason that caused that failure? There are a lot of other aspects that go into trustworthiness. like privacy, security, etc. But then these would be things that we will be working on as follow-ups. But I think we start with these three main pillars at Themis.

Shivani Sopory:

I think when most companies, when they buy services or they're implementing these machine learning tools, naively we are all thinking, well it's a system so by nature it should be unbiased. And I think what your research and probably everything you've looked at so far would indicate that's not true.

Elaheh Ahmadi:

Yeah.

Shivani Sopory:

So I guess how would you think about that or how would you articulate it to a board of directors or other kinds of companies as far as what risk there is actually in adopting some of these models and not testing them or not fully understanding their capabilities?

Elaheh Ahmadi:

The reason that comes in is a lot of, I think now the industry's more aware of it, but early on, the days of machine learning we were using... Our model is 99 percent accurate and it's the best model out there, so you should be using it and as a high-level person you'll be like, "Okay, 99 percent is pretty good, so let's use it." But the issue with the way that this [inaudible 00:08:36] measurement is coming in is that it's not taking into account these rare events or underrepresentation group or the things that don't happen, but when they happen you need to be

aware of it and that's when a lot of companies had to pay huge prices and lose their public image because they didn't think it would be important. Once you have 99 percent accuracy, what is that 1 percent?

What I would advise people when they're thinking about buying solutions from third parties is require more information, a full audit of how is this model exactly working? Show me the examples that this model is failing and show me ways that I can actually avoid that failure. Show me when the model is predicting something as X, how true is that and should I just go with that? Or if it's failing, what is the price of that to me? It's kind of asking for more. Now they've got a bias audit or an audit or more extended report on the model behavior before going into, "Okay, yeah, this is good, let's use it," and then applications.

Shivani Sopory:

In your experience, do you find that most of these tools, it is possible to actually identify why a system is doing what it's doing as opposed to it being more of a black box?

Elaheh Ahmadi:

That's a great one. Very generally speaking, I would say no. So the challenge is that a lot of these models are very complex. One example of it that I think everybody's not talking about and knowing is ChatGPT. This model is huge, it costs hundreds of millions of dollars to just train it, and when the model is that huge, being able to understand it takes even more complexity to be able to do that. To simply put it, it's not that easy solution. It's like, "Oh yeah, we just do this and it's like it gives you a full report on everything and we're all good," and that's why we exist.

So we basically, the very first step that we are trying to take towards that, being able to fully understand these models, is understanding their failures and being able to avoid it, understanding the areas that they're robust, and understanding their limitations so that when they're going outside of that limitation, you can intervene, you can bring humans in the loop, you can not let the model give any output.

Shivani Sopory:

Now are there applications for what you're doing outside of machine learning alone? Because I know there's different stages of artificial intelligence and I was reading an article that was pretty much saying that we're really just at the beginning stages of AI. Can you talk a little bit about where the evolution's going and I guess if you are willing to give an opinion, are you a proponent of the direction it's headed or do you think we're going too far and pushing the AI envelope too far?

Elaheh Ahmadi:

I think we're in a very interesting path in terms of... AI is a very big label and there's a lot of different things underneath, as you also mentioned. Right now, in my opinion, I could be completely wrong, is that the challenge is in neural networks and not understanding how these neural networks are working because the models that came before that, they were to some extent explainable. This is decision trees, regression models, a lot of that. You know the decision boundaries, so it's easy for you to know the limitation of that model. But for neural networks, they're amazing because they can have very complex decision boundaries, but it's also they're scary because you don't understand what is that limitation and what that decision boundary exactly is.

Currently, that's where we are focusing. So we are focusing on making neural networks more risk aware in the sense that when it's making a decision. Not like with our technology, we can make the model to tell you some risk scores on whether there is any bias involved in this, whether there is any uncertain model uncertainty involved in this, whether there's any data uncertainty involved in this. It's taking one step further instead of just trusting the model prediction probability, now you have some more information that you can be like, "Okay, should you rely on this or should you discard it?"

Shivani Sopory:

So when you talk about neural networks, is that a model that can learn on its own as opposed to just be giving parameters? I guess for somebody like me who might be a little bit more early on in my understanding of AI, how would you explain that?

Elaheh Ahmadi:

There are different algorithms that can be made out of neural networks. There's reinforcement learning, there are language models that can use reinforcement learning, there is transformers. I'm using a super [inaudible 00:13:22] right now, but basically that's where the innovation come in. How do you train it, how do you make it be more smart? And in the simplest way that I can put it is this really interesting and capable black box that you define some sort of optimization function or your cost function, what is good and what is bad and you throw the data at it and it magically learns. I mean there is some stuff involved in the magic but magically learns something about the data and now it can actually give you outputs with really good accuracy or performance metrics that have been used. And I don't know if that's helpful.

Shivani Sopory:

Yeah, I do think that's helpful. I think it also touches on another thing about just the information being used to run some of these models and the reliability of that data. I've got to imagine that if your inputs are incorrect or flawed in some way, your output's going to be hugely flawed as well. Is that something that you're seeing in the field, too? Not just how the algorithms build, but really the information getting put into the algorithm?

Elaheh Ahmadi:

Yeah, exactly. That's one of the, again, another big challenge with using a lot of these big data is that there is, especially some of the data, if they're taking it from internet, then you're accepting the fact that you're also taking a lot of garbage information there. A lot of the labeled dataset, those are labeled mostly by human beings even today and human beings have errors. If it's a million data samples, it's not something that you can very easily quality control it. Some of them is going to just get away. That's actually one of the things that we think about and thought about in our research is how can we, with the understanding that the data is imperfect, still try to train a model that is to some extent robust to all of that imperfection. And it goes back into trying to understand your data distribution and usually, as long as the wrong information is not taking over your data set and it's not 90 percent garbage, but rather if it's the other way, it's 90 percent good and there's 10 percent garbage, with a lot of the methodologies that are currently out there are ways to identify those 10 percent and be able to either remove it from your dataset or clean your dataset.

Shivani Sopory:

Being a start-up in a very complex and deep-tech space, I've got to imagine you've run into possibly some challenges of working with larger entities or gaining that trust with customers or potential customers. How have you managed those challenges?

Elaheh Ahmadi:

It is very hard, especially for the beginning years of our companies because you don't yet have a product. There is some things there, but it's not a full-fledged product that has been tested and gone through rigorous testing. What has worked for us is using our reputation of our founding team. We have Professor Daniela Rus, who is a legend in the space. We have Dr. Alexander Amini, who's also a very well known and very respected person in the space that he's been working on. And a lot of the people that we're currently working on is coming from, they knew us before we started the company and once they realized that we're working on this, they were like, "Oh yes, please help me. This is an issue that I'm dealing with, help me fix that."

I know that's not usually, I mean it's a privilege for us and it's not something that it's always there for a lot of other start-ups, but I'm seeing the growth aimed to bigger enterprises, coming up with programs or accelerators or incubators who connect that gap between big enterprises and new technologies and new start-ups. And one of the ones that I've recently seen is Intel Ignite that Intel started its own [inaudible 00:17:29] program and it takes in new start-ups for free, not even taking any equity, takes them into rigorous programming of how to think about business and product development. And it's also for them beneficial because then if there is anything that is useful for Intel, they can take it and implement it further in their processes.

Shivani Sopory:

For you, I mean moving from a research job to a job as a CEO, how did you navigate the switch? Did you find you already had the skill set or did you have to build up a new set of skills that maybe weren't there before? How did you navigate that?

Elaheh Ahmadi:

Definitely a lot of learning and I continue to learn a lot more. I think for us, the problem that we're tackling is very technical, it's very researchy, and me having that background was very helpful to be able to switch from being an engineer, being a researcher into now I'm the CEO, I have to talk to investors, talk to companies, talk to clients, talk business, and it's a completely different language. It's learning a whole new vocabulary and learning how to translate a lot of things. One of the interesting things that I had to learn in the beginning was just because we have a technology that we're super excited about, it doesn't mean customers are going to be excited about it and the fact that customers get excited about a product that does something for them and don't care what goes underneath that product. So accepting that and figuring out how do we change from, "Oh, we have done this five years of amazing research," into, "Okay, this is the product that has all of that underneath."

So it's a lot of very interesting changes. But I want to also say thanks to my support system of coming out of MIT, MIT had a lot of good programs and mentors and investors that really helped us to think about that and think about that change of from research to business. And also I think the culture in the East Coast is also very supportive. A lot of other founders are very willing to help you, they share their network, or sit with you and just go over your deck or your product roadmap. And I think that's very valuable to have in the community.

Shivani Sopory:

Yeah, I think the biggest challenge a lot of companies do have is taking something that you know is really great, but translating that to a customer base and getting them to understand why it's as great as you think it is.

Elaheh Ahmadi:

Yeah, it's tough. But I guess that makes it interesting also.

Shivani Sopory:

Now I understand a lot of your team initially was part of the same research group when you started. How have you thought about your team expansion and adding other people to your group or maybe even upscaling some of your current team members to take on different roles within your company?

Elaheh Ahmadi:

Most of our team even today is machine learning engineers and scientists. Since we did a lot of the research in the lab, but a lot more needs to be done to take that and make it into something that industry scale is actually usable. So in that sense we're still very actively doing research, but we also realized that our team is very technical. So we have a member from Harvard Kennedy School who has that business and more AI regulation and knows that language who has been very helpful in helping us navigate and change our image and storytelling. And moving forward I think we're basically going to bring on people who have skill sets that the rest of the team don't have, on the product building side, on the marketing side, on the business side.

Shivani Sopory:

I know it's a little bit of a detour, but I did want to just ask about your thoughts on what, if anything, needs to be done from a policy standpoint and how you think about the evolution of the technology and even where the regulations are. Because certainly there does seem to be a pretty big misalignment between how quick the technology's moving and where any regulations exist.

Elaheh Ahmadi:

I think the regulations and policies are up and coming, which is really hampering. I mean it's better than nothing. It's better than that. But I think one of the biggest challenges or issues that I'm seeing is the gap of the people who are working on these policies and regulations and the people who are building the technology and the fact that there aren't that many people who like to bridge the gap. I think we're getting better at that.

Our cofounder, Professor Rus, is involved in the global AI regulation conversation. US recently put together some of the leaders in the space to help with figuring out the best way to regulate the space. So I think we're definitely moving towards that. But the issue is the policy world moves a lot slower than these technologies do. So the challenge is while we're waiting for these regulations, what should happen in the meantime and should we stop doing the research part until the policy catches up? Which I don't think that's a good solution because we haven't done that in the past for any technology that was developed. I think one of the things that on the policy and regulation side needs to happen and is more urgent is adding something, some sort of regulation or policy in the meantime before the policy matures to the point that it can actually be scaled and have more easy way to do that, all of that checking and ensuring.

Shivani Sopory:

I do appreciate that. I actually was not aware that there's these groups of individuals that are actually coming together from both sides, from the industry and from the policymaking to come up with something that would be useful in practice. So that's good to hear from an industry standpoint.

The other thing that I was curious about, I am a big supporter of women in technology and especially female CEOs and founders in the start-up space. I think we all know the numbers aren't greatest as far as how the percentage of women CEOs and founders, especially in the technology start-ups. But you're kind of unique in that not only are you a female CEO of a tech start-up, but it's also a very deep-tech start-up, so it's highly technical. I don't necessarily want to go down the path of talking about how we ended up where we're ending up, but I was curious from your perspective on what can we do to move the needle forward? What can we do to increase the number of women starting companies as well as really increasing the number of women in more technical spaces like yourself?

Elaheh Ahmadi:

I think one of the very easy things that we can do is encouragement. I mean the numbers are definitely hugely imbalanced, but there are females in the deep-tech toward doing research and bringing that resource of you can start a company if you want to. There are all these, so it's available to you if you do want to choose

that path. So I think one of them comes from encouragement from the universities or wherever they're doing the research to let them know that this is a path that they could be taking.

Another one that I've really appreciated is the other female CEOs and founders and how supportive and helpful they've been towards me and they've been willing to do that for free. Basically just taking time out of their calendar, which is super valuable since no founder has any time and just helping, guiding me, and/or sharing their network. I think that has been enormously valuable, so we should do that more, not only female founders, but also other founders to support this minority groups and support them to grow, given that all odds are against them.

In terms of VC funding, we can look at the [inaudible 00:25:29] Center, super bad, only 1.9 percent of the VC funding last year went to all-female founders. Another part was if the female founders were paired with a male founder, that number jumped into 16 percent. So from the statistics we see that the odds are all against them so as supportive as we can be and helping with introducing them to other VCs, to other customers, I think that can change that number hopefully in the future.

Shivani Sopory:

And just curious, did you actually just reach out to female founders or other founders that you wanted some insights from, or did you get references from some of your investors? I guess how did you build your network or start to build your network?

Elaheh Ahmadi:

Mostly from the VC networks. Right now, most of the people that I know are other portfolio companies of the VCs that have invested in us. And I really love that about our VCs because they not only said that they value diversity, but they actually are doing it and they are investing in diverse founders and people who other VCs would be heavily biased against. So I think that has allowed me to get access to a really good diverse group of people. But these kind of VCs exist so you can talk to them and they would be a lot more understanding and more supportive.

Also, my MIT network has been very helpful, but I know that's not something that everybody has access to. But a lot of these networking events, especially if you're in Boston area or Silicon Valley, I'm sure there's a lot of other cities, there are a lot of networking events for founders that you can just go and meet other founders who are working in your space.

Shivani Sopory:

The other question I've got for you is currently in the environment that we're in, obviously tech did take a big hit from a public company standpoint and the public market. We have also been seeing just a slower run of fundraising that's able to happen for early- to late-stage start-ups. So I guess having your company been around for a couple of years now, have you seen a shift in how you need to be running your business in the current environment? Or I guess where you're sitting right now, is it buckle down, you've got another year or two? How are you viewing the current market and the macroeconomic effects on how you're running your business?

Elaheh Ahmadi:

Yeah, it's been very challenging in the sense that the numbers, if you talk to, and this has happened too, I would talk to founders who started in 2020 and I would ask them, "How did you raise and what did you show in your deck when you were talking to VCs?" And a lot of them would be like, "Oh, for a \$5 million seed round, I just had my team and an idea." And I'll be, "But that's not what VCs want right now. They're like, 'Where's the product? Where's the customer? Where's the revenue?'" I'm like, "It's [inaudible 00:28:16]. Come on." So I think the VC's expectations have definitely gone a lot higher on what you need to bring on the table before they even start touching their pocket to think about writing your check. So I think in that sense

it's been really tough, but also I think it's like an overcorrection to the industry and limits a lot of the people. The money should have been allocated to other companies who had better technology or better product or better strategy. It's unfortunate that the overcorrection is happening now and it's affecting a lot of founders who don't deserve to have to struggle this much. But in software, I think it's something that the market needs until we go into a more reasonable state.

Shivani Sopory:

Yeah, it's the pendulum swung one way too far. Now it's swinging the other way. We haven't quite gotten to an equilibrium, it sounds like. Well, I think we're almost at time. I do just want to thank you so much. I learned so much about not only how to think about AI, but also how to think about bias in the way we look at different kinds of algorithms, a lot of what Themis does, and just how to think about running a business in this timeframe. So thank you so much for the time, Elaheh, really appreciate it.

Elaheh Ahmadi:

Of course. Thank you for having me. It was great chatting with you and it's a pleasure to be here.

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

kpmg.com/socialmedia



The information contained herein is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavor to provide accurate and timely information, there can be no guarantee that such information is accurate as of the date it is received or that it will continue to be accurate in the future. No one should act upon such information without appropriate professional advice after a thorough examination of the particular situation.

© 2023 KPMG LLP, a Delaware limited liability partnership and a member firm of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. Printed in the U.S.A. The KPMG name and logo are trademarks used under license by the independent member firms of the KPMG global organization. USCS000787