AI and You

Transcript

Guest: Rob May, part 2

Episode 45

First Aired: Monday, April 26, 2021

Welcome to episode 45! In today's show we will conclude the interview with Rob May. Rob is a general partner at PJC, an early-stage venture capital firm focused on investing in, supporting, and building relationships with entrepreneurs who are creating the future. His expertise is in the areas of AI, hardware, B2B, and Software-as-a-Service. Before joining PJC, Rob was the co-founder and CEO of Talla, an AI and automation platform, transforming the way businesses deliver customer support. As an angel investor, he made over 70 early-stage AI investments, and I strongly recommend his newsletter on artificial intelligence, called InsideAI.

Last week we talked about how Rob got into the field of venture capital, because he started out as an engineer designing chips, and not a lot of engineers end up as venture capitalists, so I was particularly interested to find out how that happened to Rob. We also talked about brain-machine interfaces, which I was a little surprised to find out are already on Rob's radar. And we talked some about digital assistants, like the ones he built at Talla. We'll talk more about those in the second half of the interview, connecting them with affective computing, and we'll talk about what AI startups should and shouldn't be doing. Let's get back to Rob May.

I can hear that your mind is tuned to range over these huge areas, talking there about social safeguards and protection against six sigma events. And reminds me that we were talking before we started the recording here about the interdisciplinary nature of artificial intelligence and how that attracted you. In the work that you do with start-ups in AI, we're talking about something that's been characterized as "the new electricity." And so there are two ways that that could show up, it could be AI as a service: a company says, "We're providing AI to healthcare, accounting, military," the same tool, different applications, or is it more applied AI? Does one of those interest you more than the other?

Well, I'm a pragmatist at heart. And so I like debating a lot of the deep questions about the world, and I do think that's fun. Ultimately, I like seeing things come to life a little bit more, and so I think I'm much more on the applied AI side. I read a lot of the philosophy of AI and cognitive science, philosophy of mind, all of all those. I'm a big fan of, you know, Daniel Dennett and Doug Hofstadter and all these people that we've all probably read books. But I really like building and making things and seeing them come to life, and I like investing at that crossover point. And so I follow a lot of what's going on theoretically, deep, big questions, but I'm always looking for what's next to crossover and become reality. And that's sort of where I live because I'm not, you know, when I was a CEO, I'm a CEO of two companies, but I'm not I wouldn't say I'm highly operational. I'm a more strategic, sort of integrated thinker. I was more business development and sales-focused as a CEO. I'm not the guy that wants to figure out how

the logistics work for whatever thing we're doing, and so I think the applied AI is probably where I spend a lot more time.

Okay, and so talking about things that are getting ready to cross over, and we were talking earlier about chatbots and natural language processing. What about affective computing? Is that ready to take off, emotional interpretation and expression?

Yeah, I don't know. I know Rana el-Kaliouby from Affectiva, she is a good friend of mine. They've had a lot of success with it, and are using it a lot in situations related to driving and telling if somebody is tired or angry, or maybe shouldn't be driving. But I also think it depends on how far you want to take it. There's been some studies that show that maybe our emotions are always as accurately read as we think. And you have this interaction where people start to adapt to things around them, right? So you look at cell phones over the last 20 years, and people would have them, but not keep on them all the time. And now we do everything on them, right? And we're always hunched over and we're typing with our thumbs. And so I think you risk these scenarios where affective computing becomes more prevalent, people start to fake their emotions in ways to tell the computer something. Maybe, for example, a computer learns that if you seem to be getting mad, it should give you a refund. So people fake getting mad when they learn that so they'll get more refunds, right? It's like the voice trees now. They put these voice phone trees in place so that they'll be more efficient. You can find what you want. And what a lot of us do, we get on and we just start pressing zero, right? Zero, like, give me a human. I know, my problem needs a human. So I think it's definitely ready for some early applications, but it would be interesting to see how successful it is when it's rolled out at scale.

I know there's a whole site called gethuman, which tells you the fastest way to get past the phone tree to a human for all different kinds of companies.

Yeah, it's run out of Boston here where I am.

I've used it quite a few times. In order for consumer-facing AI to be more valuable, we wanted to understand it more and a lot of frustration comes when it doesn't have the information to do that. Like, say we're just having a moment with one of our children, and the teachable moment things are intense and Alexa perks up to say a package has arrived, well, she thinks she heard her name and yada, yada, yada, where you're just like, "Stop it, this is not the time".

Yeah.

And it was not reading our emotions, it hasn't been told to do that. But a lot of what we would want to enable that kind of Know Your Customer on steroids through AI would require data gathering that we also say, "Well, you can't know that much about me". Every time we find out how much Facebook or Apple or Google knows about us, all have it rubbed in our faces. We're like, "Oh my God, let's march on Congress". Is that a paradox?

Yeah, I think so. I mean, I see so many start-ups that are trying to raise money that have these ideas that like, you know, people should take back control of their data, and they should care more about privacy. But the problem is to improve your privacy and to do that kind of stuff, you

have to introduce a little bit of friction, a lot of times, into the system. If you want to increase security and privacy, you can turn on two-factor authentication for all your applications. Well, that means you got to get a text message every time you log into something, and it's going to take you eight to ten seconds longer, and people don't want that. They don't even want to log out. They just want to stay logged in, which is totally insecure. And so I think our desire for privacy and safety and security is a little more cognitive, and our desire for convenience is a little more emotional. And so I think the desire for convenience keeps winning. But that's also part of the reason that I think some of this is going to have to be legislated; because there are certain situations where core human emotions can be manipulated, you know, gambling, fast food, these kinds of things, where, if the government doesn't step in, I mean, yes, the market will make it a race to the bottom, right? And so, while in general, I'm not in favor of heavy government regulation of technology. I do think this is a place where, I think Europe has really led the way with GDPR and some of the stuff - the right to be forgotten - and some of their laws, and I think I'd like to see us do more of that in the United States. And I'd like to see more movement towards worldwide treaties about how we treat data, personal privacy data, cross border, and all that kind of stuff, because it's really, it's really complicated right now.

Are there areas where you see start-ups wanting to go, that your advice would be in general, nope, don't go there, for whatever reason?

Oh, definitely. There a handful of problems that people always want to solve that sound like good ideas that never work because of other structures. So I can't tell you how many companies I've seen that have done this, like, "Well, look, here's how we'll fix advertising online. It's micropayments for content". So you want to read this article, you pay half of a penny. You can just do that, and those things never seem to work because you have a scale problem, right? You have to get them out everywhere before they're valuable. And then people want transparency. I've seen so many of these sites, Secret was one of these where you could post stuff, you could tell truths about people. I've seen sites where people like, you know, honest, real feedback, but anonymity always breaks these systems and they end up like, people don't really, like it hurts to get real feedback a lot of times. Like society is kind of built on lying to each other, to some extent to maintain social ties. And so there are some things like that own your own data, privacy, have a little bucket that you share with Facebook when you want to be advertised to. is one. Yeah, so we see a lot of these, and then those are very specific. They also come in conceptual flavors, like super technical people think about a lot of things that they think other people would like, that other people wouldn't like, and sort of try to build some of those companies as well. So yeah, we're always dealing with kind of that those paradoxes of what people think should exist, human tendencies that are pitted against each other, and what's ultimately really scalable and possible as a business model.

What have people in start-ups not addressed that you would love to see someone get into?

Well, let me take that from a problem space and a technical space. I think from a technical space, I would love to see more companies working on AI that's not neural network-driven. So I've seen this in academia, which is, you can look at Bayesian analysis, symbolic logic processing, evolutionary algorithms, there are other forms of AI. And I'm not aware of many companies that

are using those in sort of production systems for things. We're starting to get there. I've seen some work at IBM, for example, on using neural networks with symbolic logic; maybe use neural networks to generate symbolic logic. And then symbolic logic makes it easier to do logic and reasoning and higher-order functions that neural networks can't currently do. I think that's interesting. I'd like to see more people working on that tech. I've seen a few but there's not a lot of start-ups doing it. And then from a problem, space perspective, I think there's a bunch of interesting problems to solve. One that may or may not end up with an AI solution is I think b2b buying is broken, right? So when you have a new type of software today and you want to sell it, you make a lot of content, you put it online, but you do a lot of cold emails and cold calling, and everything else, and you spam people, and it's annoying. But people do need to be aware of products and they do need to buy them. I used to go to this buyers' conference, I think it was called the CIO forum or something like that. And, you know, they would invite 150 sort of CIOs, and they would pay for them to come; and they would say, "Look, you're going to hear from your peers, we're going to have a fun poker night, you're going to play golf. And you're going to listen to eight pitches, half-hour pitches from companies". And you would ask the CIOs, these are heads of companies and cities and state governments and things like that, it's like, "Why did you come to this thing for two days to get pitched", and they're like, "Well, I have a good time, and I get to see a bunch of people I don't normally see. And I also don't mind the pitches, like I need to buy stuff, and I want to know what's out there, and it's better than getting 30 phone calls a day from random people". And so it seems like there should be some way for somebody to get in and fix that, like, know what you actually need and screen stuff, so that you only actually get pitched and only take pitches on stuff you're really in the market to buy or would really be useful for your company. I mean, I started two SaaS companies, and we would get pitched on things like equipment leasing, like, we don't lease any equipment, like, what would I use this for? So I think that's an interesting problem; then with respect to AI, I think most of the most interesting issues are around edge computing right now; because people are predicting that this will be bigger than cloud computing. You're going to put machine learning models in every device. Every camera is going to have a facial recognition model or something else, things are going to have voice recognition in them. And as these models get better and more accurate, they do it by getting bigger, typically. So how do you compress these models that run on these large server-side chips and put them on very small, cheap microcontrollers? Because it's not, like if I have a camera that is connected outside of my house, and it's trying to look for the faces of certain people and identify if it's my family member or not, right, should I let them in or something like that. I can't just, it's not cost-effective to send that up five times a second backup to the clouds for inference, and whatever, is there a person, and they're not like, you have to put that inference down on the camera, it's more economically feasible. But putting that model to the camera is hard, depending on the hardware that you're using, without driving the cost of the camera off, right? I mean, I'm not going to pay $700 for a video camera that previously I could get for $89 mounted on the front of my house. So I really have been looking at a lot of companies that are working on how to either build those models with fewer data by pruning the data sets or compressing the models on the other side, so that they can fit on these things, or just different tools like that. So I really think the electronic design tools around neural network models going into production, a big place with a lot of problems to fix.

That's interesting. I think the conventional wisdom would be that 5G plus Starlink, and eventual 6G with terabit speed, would say that you could go to the cloud cheaply.

Yeah, well, you're still going to have issues. You're going to have latency for some issues. If you're talking, you know, 20 milliseconds, 100 milliseconds at an autonomous vehicle that might matter, right? And I think you're still going to have, you know, cost issues and privacy issues, internet connectivity issues, there's going to be a lot of reasons to push this stuff to the edge.

True. Are we in an AI bubble? Do you think at some point, this is going to pop the way of the dot.com, and we'll be going wow, that was like bell-bottoms or cat memes? We were a little carried away with ourselves there. Do you have had companies come to you that have got their AI label on what they're doing, and you look at you go, "Dude, that's just linear regression, you can't call that AI"?

Yeah, so I'll say yes, and no. I think we're in an AI bubble, in the sense that there are a lot of things being pitched as AI that aren't AI. As intelligence becomes part of everything, I think there it'll just become the norm. It's like, nobody says, like, "We're an internet company" anymore; because every company has to be on the internet. I do think we've over-invested, particularly a lot of the tools and platforms and things like that from an AI perspective, but that over-investment's good, right? It lays the groundwork. I mean, it's a lot of speculation, but it lays the groundwork for the real use cases that come later. But I don't think it's necessarily bubbly in the sense that it's going to pop and go away; because you've seen a lot of real use cases and real improvements, particularly in the big tech companies, right? What Tesla's done, what Amazon's done, what Google's done, Microsoft, Apple, even places you hear less about like Salesforce and Oracle are doing a lot with AI. It's making a lot of things better, and it's really rolling out; and so AI engineers aren't going away, machine learning engineers. The technologies are going to keep moving forward, but yeah, there are too many companies in some spaces, the MLOps space is a perfect example where everybody saw the rise of DevOps and some of these companies, they got big, you know, Splunk, and Datadog, and people and everybody said, "Wow, we want to have the MLOps version of this, where instead of looking at all the server, and compute stuff, we're looking at the data and managing the data flows and everything around the models. And there are so many companies in that space, we haven't at PJC, we haven't made a bet, because I don't know how you tell who's going to win. I think it's going to be a little bit of randomness and luck, and who raises the most money first, and so there are definitely some bubbly sub-genres of AI for sure.

So you've got this newsletter, Inside AI with a very large base of subscribers, very successful, that's how I first learned about you. What was your goal in establishing that? What has it done for you?

Yeah, so I started that in 2015, because my first company was a cloud security and backup company, and I needed to rebrand myself as an AI guy if I was going to start an AI company and do AI investing, which is most of my angel investing. So before I joined PJC, I made 74 angel investments, and probably 65 of them were more AI-focused. So if I was going to do that, I was thinking about how do I get deal flow? And how do I rebrand myself? So in 2015, I

started writing the newsletter was originally called Technically Sentient. And, you know, when it got to about two or three thousand subscribers, Jason Calacanis, who's a friend of mine was doing the inside dot com thing, and he said, "Hey, we want to launch an AI newsletter. I love yours, like, let's just put it up here, you can have full editorial control". So I put it there, and I was writing every other week, and then I started writing weekly commentary, and we changed the name to Inside AI. They grew it to 10,000 subscribers in just a matter of weeks, right?. And then now it's got about 30,000 subscribers. It's one of their, I think it's the newsletter that we have a couple of 100 newsletters, and mine has the highest open rate. And then we added, you know, about a year and a half ago, we added somebody who would write the news daily, and I still just sort of focus on the weekend commentary. So yeah, it's, it's cool. That's why I did it. It's very scalable for me to send that, you know, most of the people that I personally know who read it, they feel like they hear from me all the time. And so in a time like COVID, when I can't have coffee, our relationship maybe feels closer from their perspective, when I reach out to say hi, then I otherwise would. And I also recently have started writing on Substack at investinginAI.substack.com; and that's where I'm going to write more about the just the investing side of some of this stuff, you know, business models and markets and things like that. So I like to write mainly just to, it helps me with what I think I mean, I really do it for me more than the readers.

I know what you mean, I wasn't originally planning to write a book. When I started this, I was going to go to videos and and talks, but I was having so much trouble ordering stuff that I thought, well, I'm going to write this down so that I'm forced to figure out what order to explain it in, and then I thought, "Oh, this really should go out as a book".

Yeah.

How should people keep up with what you're doing, and find more about your different ventures and publications?

Yeah, so the best thing would be, you can follow me on Twitter. I'm just Rob May, or investinginAI.substack.com is probably the place where you'll see more stuff about me personally and some of that information.

Any advice for someone starting out in your field, if you're mentoring them?

If you want to get into venture capital, it's really hard, right? You either have to be a successful CEO a couple of times, or a successful entrepreneur, or you have to work your way up a very long ladder from you know, associate analyst principal to a partner. So I would just encourage people, whatever field you're in, try to stand out, do something different. You know, write blogs, podcasts, organize meetups, when COVID's is over, like, whatever you can do to just rise above the rest a little bit, and it's hard. I wrote Inside AI, when it got started. I mean, six months into it, I had like, 150 people read it right? And it seems very depressing, but a lot of this is just keeping at these things, right? Consistency and discipline, and all that takes you a long way.

How would you like AI to change your life in 10 years?

Well, I would really love something that recorded everything I did all day. And I could ask questions like, "Hey, when I was talking to my friend Steve the other day, and he mentioned that restaurant, and that made me think about that dish that I had, like, what was that"? Those are the kinds of questions that I can never remember the answer to. I'd like to have something like that. I'd love to see AI revamp education. I think people learn so differently, and the educational system has been so homogenous that I think, you know, to be able to tailor education to every child so that they can learn at the appropriate sort of pace and the things that are, make things interesting to them. I think could really, really change the world.

That's a great place to end this. Rob May, thanks for coming on the show.

Yeah, thanks for having me, this is fun.

That's the end of the interview. I hope you're having as much fun with these interviews as I am, because I find our guests utterly fascinating, and I wish I could talk with them for much longer than the hour we typically get to spend together. Maybe one day we'll be able to make something like that happen, like say, a panel discussion, or an all-day workshop, or a live lecture series. Which reminds me, that if you want to see anything like that happen, it all depends on listeners. So this is me telling you to share the show and tell your friends about it and give us a five-star rating on iTunes or wherever you're listening to us. I suck at marketing, so I don't say this nearly as often as a lot of podcast hosts do, but it really does make a difference. Our listener base continues to grow with every show and that's thanks to your sharing about us, which all helps to win the battle for attention.

In today's news ripped from the headlines about AI, the National Highway Traffic Safety Administration (NHTSA) announced that it has opened 27 investigations into crashes of Tesla cars. NHTSA said in July that its "(Special Crash Investigations team) has looked into 19 crashes involving Tesla vehicles where it was believed some form of advanced driver assistance system was engaged at the time of the incident." In one example, Michigan State Police said a parked patrol car was struck by a Tesla apparently in Autopilot mode while investigating a traffic crash near Lansing on Interstate-96. No one was injured and the 22-year-old Tesla driver was issued traffic citations.

So this is important, and it's gonna be hard to parse out what's important and what isn't, because every incident involving autonomous vehicles is going to get a huge amount of attention. An AV runs over a cat it's gonna be headline news. An AV ran over a person outside Tempe, Arizona and Uber stopped that testing. Now, there are 30,000 people killed every year by traffic accidents in the United States alone and they don't get hardly any attention so are we being unfair to the autonomous vehicles? Well, certainly they're going to get more than their fair share of attention because they're new and we don't know what they're going to do. On the other hand, they're driving a lot less than human drivers are right now, so really we have to look at statistics to find out whether the attention that they're getting for these crashes is really warranted. Elon Musk said that Teslas are already safer than human drivers, but we need some kind of analysis to tell us whether that's true, and how true it is. The entire point of making autonomous vehicles is to reduce that 30,000 a year figure by at least 90%. Otherwise it's just not worth the billions of dolalrs being spent by that and many other companies. This may create a Catch-22 where the attention that those vehicles get from NHTSA and especially the media may defeat the introduction of them on a wide scale even after they become generally safer than human drivers. But I'm thinking that what may unwedge that is insurance companies. Because they deal in data. You

can bet that they already know just how much safer or less safer AVs are than human drivers. When that equation comes out on the side that works for their underwriter,s they will push this forward because they stand to make money off the widespread use of autonomous vehicles as soon as they are significantly safer than human drivers.

Next week, I'll be talking with Rajiv Malhotra, who has a new book, *Artificial Intelligence and the Future of Power: 5 Battlegrounds*, which is particularly focused on the impact of AI on the economy and social framework of developing countries like India. That's next week on *AI and You.*

Until then, remember: no matter how much computers learn how to do, it's how we come together as *humans* that matters.

http://aiandyou.net