SDS PODCAST EPISODE 895: THE FUTURE OF ENTERPRISE AI: INVESTOR SHAUN JOHNSON REVEALS WHAT ACTUALLY WORKS

Jon Krohn:	00:00:00	This is episode number 895 with Shaun Johnson, co-founder and general partner at AIX Ventures. Today's episode is brought to you by the Dell AI Factory with NVIDIA and by Adverity, the conversational analytics platform.
	00:00:20	Welcome to the SuperDataScience Podcast, the most listened to podcast in the data science industry. Each week we bring you fun and inspiring people and ideas, exploring the cutting edge of machine learning, AI, and related technologies that are transforming our world for the better. I'm your host, Jon Krohn. Thanks for joining me today. And now, let's make the complex simple.
	00:00:54	Welcome back to the SuperDataScience Podcast. We've got a terrific episode today with an iconic trailblazer AI investor, Shaun Johnson. Shaun is co-founder and general partner at AIX Ventures in San Francisco, where he's led deals into companies including Perplexity, Chroma, and Workhelix. He's a former VP of engineering product and design at Lilt and a former VP of product and design at NimbleRx. He holds a master's in electrical engineering from Stanford and an MBA from Berkeley.
	00:01:25	Today's episode is well suited to any listener to this podcast. In it, Shaun details how having investment partners like Richard Socher and Christopher Manning who are practitioners actively building at the cutting edge of AI gives AIX Ventures an edge. He talks about what it takes to become one of the few thousand people in the world pushing the AI frontier, the surprising strategy that makes enterprise AI adoption 10 times easier, why some AI startups are better off building in red oceans full of competition rather than seeking blue ocean opportunities, and the reason big tech companies are buying AI talent without acquiring the actual startups. All right. You ready for this excellent episode? Let's go.

	00:02:10	Shaun, welcome to the SuperDataScience Podcast. It is a treat to have you on the show. Where in the world are you calling in from today?
Shaun Johnson:	00:02:19	Shaun, thanks for having me. I'm calling from San Francisco. We're in our offices here in South Park.
Jon Krohn:	00:02:23	Where better for an AI investment firm to be than San Francisco. I'd expect you nowhere else. So you're a founding partner of AIX Ventures, which is an AI-focused, early-stage investment fund being led by active AI practitioners. So, let's parse that for our listeners. What are the advantages of being an investment fund that is first, AI-focused, second, early stage, and third, led by active AI practitioners?
Shaun Johnson:	00:02:53	Yeah, it's a great question to kick this off. I think, first, being AI-focused, every investor out there, our job is to find alpha for our own investors, right? How do we think about returning, outperforming funds? We just have the strong belief that we're right in the middle of an inflection point, right? We saw transformers, ChatGPT, and machine intelligence is growing every day. We think the next dozens of $100 billion, $50 billion companies are coming from AI, and so this is where we're heads down. We think this is transformative for industry, for humanity. Other investors could find alpha doing non-AI, but here we are I think is number one. Then two was the operator. The operator.
Jon Krohn:	00:03:59	We can do that. I had early stage as number two.
Shaun Johnson:	00:04:01	Early stage. Yeah. On the early stage side, look, we do two things. We think about backing entrepreneurs at the earliest stage because that's how AIX started organically, right? My co-founder, Richard, has been investing in the space since 2016 and Chris Manning long before that. That's just early stages meeting founders where they are

when they're starting to think about these ideas and then you bring that into your first fund. Now, also, I should say that our own investors think about early stage versus growth stage. Growth stage are lower risk, lower returns and also more correlated to the public market and early stage are not, right? They're less correlated to the public market and then also higher alpha, higher returns, and so that's quite a nice product market fit for us too.

Jon Krohn: 00:05:09 Perfect. Yeah, it makes a lot of sense. I was talking to you about this briefly before we got on air, but Richard Socher has been somebody that's been an iconic person in AI in the world for me, I guess, as an AI practitioner. It's going back roughly a decade that I've been studying and teaching from his deep learning for NLP lectures at Stanford in Chris Manning's class who you also just mentioned there, who's still a Stanford professor. So these active AI practitioners that you have, that you've co-founded AIX with, I mean, it's an amazing set of people that it's hard to imagine people that I look up to more. Yeah. What's the advantage of working with AI practitioners as opposed to, say, career investors?

Shaun Johnson: 00:06:00 Yeah, we think you need both. The most important thing I think for any venture firm is to orient yourself around the founder, and so we think founders need... They will for a long time now be building at the intersection of AI plus a sector they want to disrupt. So, how do you support that founder? Well, we think it's hard to support on the AI side being a full-time VC, right? We think you're better off being out AI is moving so fast, be out practicing at the cutting edge of AI. We mean the cutting, cutting edge, like the Richard Socher, Chris Manning, Anthony Goldbloom cutting, cutting edge.

00:06:39 Then when we think about on the sector side, we think about full-time VCs, market-focused traditional VCs. You have to bring those two together and I think the way we

brought it together is what's quite unique, which is in a very engaged way, right? So Richard and Chris and Anthony are in every single one of our investment committee meetings. I mean, Anthony's working in our office just a couple feet from me now. So it's just this very engaged model where founders walk in and see that, "Great. They have both sides of this equation and it's differentiating for them."

Jon Krohn: 00:07:16 Nice. Well explained. You were talking earlier about how investors, you have this thesis that we're at this inflection point in AI and you're preaching to the choir a bit with me and probably our audience as well. I think we're on board in that belief. It seems obvious with how dramatically rapidly the cost of compute around the same level of capability advances. I was just at the time of recording the episode that I published today of this podcast, I had some stats from the Stanford state of AI report that I was talking about in the episode. If you take something like a 60% accuracy score on MMLU, it took hundreds of billions of model parameters in something like Google's PaLM just a couple of years ago to be able to get to that level of capability. Now, today, we have models that are 200ths, one-200th of the size, just a few billion model parameters that are at the same capability and the cost of compute has gone down even more.

00:08:29 So depending on exactly the kind of task, it can be several hundred times cheaper to 900 times cheaper per token to be getting the same level of compute. So say you fix something at GPT 3.5 level capability and you track the cost of that over the past couple of years, it's up to a 900x reduction in cost. So with those kinds of tailwinds, it seems obvious that there's a lot of opportunity in AI, but there is also... I think there has been overhype recently as well. Yeah. What do you think about that? In terms of Gartner hype cycle, maybe we're now in a trough of disillusionment, which I could get into more detail on this

later, but it seems to me like that actually means there's a lot of opportunity for real-world implementation coming up next.

Shaun Johnson: 00:09:20 Yeah, I think with any great technological inflection point comes a hype period too, right? So right now, I think you're seeing that manifest itself in a couple of different ways. One is the efficacy of AI in consumer's lives or in the enterprise, right? Then another way is you're seeing an abundance of capital pour into the market looking to support these companies. Some companies have sprung up that maybe there's no need for, but there's so much capital. Entrepreneurs will give it a go and you're going to see softness and returns because of that. So I think it's both true. I think even though maybe broadly there's hype or we're in a hype in that trough of disillusionment, it's too general, right? I think you have to look at specific applications and ask where is that in the cycle, right? Because we see applications today regularly, founders coming in and talking to us where there's quite a blue ocean, right? There's a lot of opportunity. It's not all red. Not everyone's competing. Maybe they have an angle on it that's quite different, and so it sort of just depends.

Jon Krohn: 00:10:43 Right. That makes a lot of sense. When you're evaluating these early stage AI startups that you invest in, what are some of the non-obvious signs of product market fit that you look for? Or I guess even more generally, what are you looking for in those investments? How do you evaluate that the market is ripe for that particular type of AI in that particular application?

Shaun Johnson: 00:11:07 You know what, early stage is really a game of people. You back the founder or founders when they're just having their idea, maybe they have some prototype or some product, but it's really a people bet. They will go out there with their vision to change the world and they will learn quite a lot, and that will result in pivots, micro and

macro, and so we can't say... I don't think VCs are genius market timers, right? I think they can have a sense of that, but then also recognize that the founder will do what they need to do. And we really just look at investing in teams that can execute at the speed of light and pivot however much they need to find a resonant point between what their offering is in the market and then get to that great growth trajectory.

Jon Krohn: 00:12:14 That makes a lot of sense, but then it begs a similar kind of question, which is how do you then identify that kind of founder or that kind of founding team? I've had guests, I've had investors on in the past who have said that with AI startups, they typically look for this kind of three-legged stool of A CEO, which is somebody who's great at selling the idea, a CTO who obviously is highly technical, but then in an AI startup you also have this AI expert, where the CTO is maybe more concerned with platform scalability, reliability, those kinds of concerns. You have this third co-founder that is the AI expert at or near the cutting edge like you described Richard Socher or Chris Manning might be with their research. Does that ring true to you as well in the teams that you are investing in?

Shaun Johnson: 00:13:13 Not that I would say 100%. The way we think about it, we start by looking at the team and assessing two factors. One is AI-nativeness, right? Do we consider this team to be quite deep in AI or not? Then market-savviness or commercial-savviness. Do they have expertise in this area? Do they have any right building into this market? That's really where we focus. Then we ask ourselves, given a team and given where we think that market landscape is, where will they need to improve? It's never perfect. You don't find teams that are always optimally AI native and optimally commercial-savvy. So if you invest in a team that's more AI native and less commercial-savvy, then the question is how do you de-risk the

commercial-savviness with the team? Maybe that's advisors, et cetera.

00:14:08 Then I think when technological inflection points like we've seen with ChatGPT happen, what happens in the market is you have a number of consensus applications that are now possible, right? Everyone knows that we should do AI-powered tutoring, and so everybody is like, "Well, let's build AI-powered tutoring." But what you need to do there we think is invest in extreme AI-native teams that can actually bring experiences to consumers that other teams just cannot. As the years go on and you start getting outside of this consensus-driven investing, you go back into the market-savvy investing where you don't need as much the AI native teams. They still be very good, but it becomes more important to have a market insight that is non-consensus.

00:15:11 So the way we think about it, if you think about SaaS investing, let's call it five years ago, SaaS investing, there wasn't a ton of differentiation on the MCV stack. It's not like you're like, "Oh, the model technology is unique. The database is unique, or the controller is unique or the view is unique." It's all Mongo and MySQL and it's in the middle where it's, let's call it, Node.js or Ruby and then it's React or HTML, CSS, JS, and that's all commodity, right? It's just like, "Well, what's the idea? How are you going to configure?" AI will get there within the current framework of the technology. Now if we have a new architecture come out that does replace transformers, then a game on again. Right now, a whole new set of consensus that will be made against what that technology can create. But right now I think what we're transitioning from you really need AI native teams in a consensus world to you're going to start needing more market-savvy teams in a non-consensus world.

Jon Krohn: 00:16:29 This episode of SuperDataScience is brought to you by the Dell AI Factory with NVIDIA, two trusted technology leaders united to deliver a comprehensive and secure AI solution. Dell Technologies and NVIDIA can help you leverage AI to drive innovation and achieve your business goals. The Dell AI Factory with NVIDIA is the industry's first and only end-to-end enterprise AI solution, designed to speed AI adoption by delivering integrated Dell and NVIDIA capabilities to accelerate your AI-powered use cases, integrate your data and workflows, and enable you to design your own AI journey for repeatable, scalable outcomes. Learn more at www.Dell.com/superdatascience. That's Dell.com/superdatascience.

00:17:17 Right. That sounds like a great balance. I guess I was being oversimplistic in my thinking of, "Yeah, this is the founding team." It makes perfect sense that obviously every situation is different.

Shaun Johnson: 00:17:30 The last thing I'd add to that question, Jon, is this comment on do you need the AI expert on the team? So you have your CEO that's market-savvy and you have your CTO that's builder, and then do you need this AI expert? I've seen this on lots of teams, even my last team at Lilt. I think what the best way to orient a team is to have AI engineers, folks that can build very adeptly with the technology in production that are also savvy enough to be reading the papers and understanding how the technology is changing and be able to integrate that into the stack. I don't think you need a PhD that can just sit there and read papers. Ideally, your folks building in production are also savvy enough to read papers is our take on it. Good news too, because AI PhDs are expensive.

Jon Krohn: 00:18:30 Yeah. Yeah. Also, good news is it's getting easier and easier to read AI papers because you can use your

favorite LLM to help you understand what the heck's going on and get an explanation for the pertinent parts. There's the LLM tooling obviously a great opportunity for startup founders in terms of the products that they can be building, and therefore, also for the investors that are investing in them, but even just for any of our listeners who want to be making the most of understanding the great information that's out there, the exponentially more AI papers that are being published all the time, it's easier and easier to understand them and it's easier and easier to code them up with tools. At the time of recording, Claude 4 has just come out and they've really focused on the ability to co-generate in there. Yeah, it's a really exciting time to be in this space.

Shaun Johnson: 00:19:30 Indeed.

Jon Krohn: 00:19:31 So when you're scaling up AIX... Recently, you just closed your second fund for $202 million, congratulations. As you scale up, as you invest in more companies, you probably need more and more AI practitioners. Another name that just came up from the research document that we prepared for you here is Pieter Abbeel who has been a guest on this show as well some years ago now, but still one of the most fascinating episodes that we've had. Pieter's an incredible person, and so that's another name associated with AIX. What kinds of mechanisms do you have in place to ensure that founder quality, as well as these kinds of practitioners like Pieter, like Richard that you have advising... You said that Richard today is still in every investor meeting. Are you concerned about that scaling or does that seem straightforward to you?

Shaun Johnson: 00:20:28 Yeah. I mean, generally, VC doesn't scale, not AIX specifically, all VC, right? It's a very heavy services business, and the reality of it is every VC out there is creating a new portfolio of anywhere between, let's call it, 15 and 40 companies every three to five years. Recently,

it's been on the earlier side of that. So how does VC work, right? You're bringing on 30 companies every three years, 10 a year approximately. How do you continue to service everybody as these portfolios grow and grow? The answer for a long time has been that maybe somewhat controversial, but it's the truth. VCs invest in companies 10 a year, and those companies usually raise for 18 to 24 months. They raise capital to build for 18 to 24 months. At the end of that period, actually not even the end, 12 months before they run out of capital, they start raising again, right?

00:21:37 Then they go and they take more capital from other venture firms too. Those venture firms start picking up the responsibility to bring these companies to the next stage. So it's kind of this founders support. Founders go through this journey of multiple VCs through the lifecycle of their company being the most prominent people on their board, and so it's always changing over. An example of this is Richard invested in Clem and team at Hugging Face when they were taking 224 at Stanford, and that was some time ago, 2017. That's a $4.5 billion company now. Clem is surrounded by the best at Sequoia and Lux and others. Clem's not calling Richard every day. He has other problems. They keep in contact and Richard still advises, but it's a much lighter lift than initially.

Jon Krohn: 00:22:45 Clem has also been on the show. He was in episode 564, I just looked up here. That's another great episode. I guess I should probably look up the Pieter Abbeel episode. That was 503.

Shaun Johnson: 00:22:56 Nice.

Jon Krohn: 00:22:59 Yeah. Going back almost four years now to that, so listeners can check out. Amazing people. It's interesting to me when you talk about Clem and Richard and their interaction like that. I don't know if this is even going to

be a controversial question, but I suspect this varies all the time. But when there's a really successful investment that you make like that, especially from such an early stage like Richard had with Clem and Hugging Face, do you think you're more incentivized to stay in touch and really be close with those really big successes, or does it end up being the case that you end up having to spend most of your time with the startups that you're like, "There's just something about the product market fit or there's something we just don't have right"? I don't know if that's a controversial question, but I'd love to-

Shaun Johnson: 00:23:55 Yeah, I think the saying is your fund is made by your winners and your reputation is made by helping your non-winners, right? So we have responsibility to all of our founders. Indeed, I'm working with founders that... I speak to Arvind regularly. Arvind's killing it at Perplexity. He doesn't need a ton of our help now, but we keep that relationship. But I'm also spending a ton of time with folks that haven't been as successful as quickly as Arvind, but still have tailwinds and they're seeing progress and they're just looking to find the right way to position their company in the market. So yeah, I think both is the answer. We have responsibilities to both, both from a fiduciary point of view and also just helping founders and a reputational point of view as well.

Jon Krohn: 00:24:57 Nice. That was a really great answer. You managed to make what I thought might be controversial into something that sounded quite ordinary. Thank you. Our research dug up an interesting question here, specifically related to climate change. So there's a bit of context here. Your newest fund, it'll invest millions of dollar checks each in applications across enterprise software, healthcare and the climate. So I want to dig into that climate one a little bit. Due to capital intensity, it sounds like you won't be investing in raw model builders, which makes sense. I mean, we're now at a point where it's at

least hundreds of millions of dollars to be building state-of-the-art AI models. But in addition to that, those models, not just at training time, but especially at inference time, when they get used a lot, these huge models, they can have a lot of climate impact. They require a lot of energy, a lot of water to cool them.

00:26:09 Now, there are some solutions out there. I know that a lot of the hyperscalers, Google, Microsoft, they try to use renewable energy at best or nuclear where they can at worst to try to limit climate impact. I know that there is a lot of innovation around water usage where the water that evaporates actually gets trapped in a system that allows it to be collected so that you're not using much net new water while running these kinds of server centers. But given that you're specifically investing in climate AI startups, I'm curious what your thesis is or what your perspective is around how is climate tech and AI become increasingly intertwined, the kinds of applications or infrastructure at this intersection that can both create outsized returns for you and your investors while also reversing environmental impact?

Shaun Johnson: 00:27:17 Yeah. Yeah, it's a good question. So when we were starting our fund two, we thought deeply about climate. There's quite a lot of work going on innovation in the climate space, but what we were quite keen to understand is the intersection of climate and AI, where we thought about novel AI solutions being applied to climate. Our team basically concluded that there actually weren't a sufficient number to invest in that area thematically at the early stage. Also, they tend to be a little bit more capital-intensive because atoms are just more expensive than bits. So we actually have not pursued that as a strategy in our fund two. Our fund two is almost squarely focused on enterprise applications, both horizontal and vertical, as well as tech bio, the intersection of AI and bio, but not quite climate just yet.

Jon Krohn: 00:28:25 I gotcha. I gotcha. I gotcha. All right. So let's move on then to other aspects of portfolio that you build, the decisions that you make. In a recent interview, you discussed two kinds of investments that you make in your portfolio. One of them is heat seeking and the other one are what you called truffle hunting bets. So, what are these? How do they complement each other as well?

Shaun Johnson: 00:28:53 Yeah, this comes from the Andreessen guys. Chris over there, I think, talks a little bit about this, which is heat seeking is a company, a deal of group of founders that the venture market knows a ton about and is very excited about and there's a lot of heats surrounding that deal, and that means they tend to be a little bit more expensive too. Then truffle hunting is just that too, which is you're trying to find the founders that are the ideas, the companies that folks just don't know, maybe it's that they're unaware of them, but also maybe it's that they're aware of them, but they don't believe that what they're saying is true. It's non-consensus, right? You want to have a portfolio of both, right? So heat seeking won't always be right. There's this adage that if you go back and you look at any vintage of fund and you track that vintage to what was hot at that time, those are unlikely to be the alpha returners, right?

00:30:07 Investing in those trends during that vintage is not what creates strong financial returns. It's actually the non-consensus bets where you can get in arguably cheaper. Then also if you're right, it could be a very, very big win. So you just have to create portfolios of these companies both heat seeking and non-consensus. It's unique when you're in the middle of this of consensus world where you just have ChatGPT November 2022 become a thing. Everything becomes consensus very quickly, right? Then you have to ask yourself, "Okay. Are we going to diversify with non-consensus or are we just leaning into this consensus world and taking a big bet

there?" Folks play it either way. We have companies that are both.

Jon Krohn: 00:31:00 Nice. It's great to get that insight into the kinds of investments that you make. The latest investment that you made, at least according to our research at the time of recording, and you can correct me on this, but we have it as Workhelix.

Shaun Johnson: 00:31:12 Not the latest but recent. Yeah.

Jon Krohn: 00:31:15 Recent. Yeah. Workhelix is a platform to understand AI's impact in a task-based step-by-step manner, which is revolutionary idea. I'm sure I'm going to butcher Erik's last name.

Shaun Johnson: 00:31:33 Brynjolfsson.

Jon Krohn: 00:31:33 Brynjolfsson. Erik Brynjolfsson, he said that this approach could help unlock a trillion-dollar opportunity. Of course, that's the kind of thing, I guess that kind of confidence is what you look for a lot of the time in founders of AI startups. But could you explain this kind of what Workhelix is doing differently, what this means to have this task-based, step-by-step approach and the kinds of enterprise cultural shifts that would be necessary to unlock the value from a platform like that?

Shaun Johnson: 00:32:04 Yeah, sure. So if you think about how enterprises are adopting AI today, how does this happen? So I think there's one of two ways. They're learning about it through the media, through what they're hearing companies that are providing solutions, customer support, the Windsurfs and Cursors of the world engineering, and then they go and they do that and they give that a go. Another way is the board. The board of these companies is learning about AI and then they're putting pressure on the C-suite to have an AI strategy and to implement AI. But the real

question is what's going on inside the organization? Every organization is made up of sub-organizations, and those sub-organizations have jobs and those jobs have tasks. Then the question is what is the task profile or characteristic of any organization? What are those jobs and those tasks? What are they? If you can go in and get access to the raw data in an enterprise, you can surface them and then you can map tasks to actually GenAI solutions, and not only tasks but bundles of tasks that cross job descriptions.

00:33:30 So what enterprises are learning is that, "Oh, we have much more labor," for example, "in manual tasks in marketing than we ever thought we had." This maps this other GenAI solution very well. So the work you seem could expose all of that and then partner with the enterprise to actually implement and most importantly measure, show that there is efficacy from this solution in a proper A/B test. When you're done with that, then you start over. You go and reassess. The organization's changed, you want to reassess. There's this really interesting thing happening in the market, which is folks, investors investing in roll-ups, this idea that you should buy companies and make them older, maybe more service-oriented companies, a lot of OpEx and roll them up and power them with an AI-native system, reduce your OpEx and potentially even increase the top line through productivity enhancements and then have a more valuable business.

00:34:49 You'd ask yourself, "Well, why would you do that? If the enterprise is adopting AI very readily, why would you focus on roll-ups?" The answer is there's friction in the enterprise, right? This is massive behavior change selling into the enterprise AI-native products. I think that's why if you look closely at some of these companies, you actually see professional services increase. I think we're seeing these companies, they'll come up with SaaS

models or usage-based models, but they also have professional services line item, and that's because of the change that is required in the enterprise.

00:35:33 There's a lot of pushback. You've seen on various social media recently CEOs be quite vocal about, "No, this is the way it's going to be." Some have put their foot down and just said, "We are adopting." I was talking to the CEO yesterday and I asked him, "Are you using AI in any real way?" He texted me back and he said, "We had to mandate it. Every Friday is AI day, where just like, 'You have to be using it. There's no objections.' " A lot of what's going on is trying to figure out how to change the behavior and that's why folks are investing in roll-ups. When you're rolling everything up and then just doing a larger riff to the organization, it just arguably is easier than that behavior change.

Jon Krohn: 00:36:21 This episode is sponsored by Adverity, an integrated data platform for connecting, managing, and using your data at scale. Imagine being able to ask your data a question, just like you would a colleague, and getting an answer instantly. No more digging through dashboards, waiting on reports, or dealing with complex BI tools. Just the insights you need - right when you need them. With Adverity's AI-powered Data Conversations, marketers will finally talk to their data in plain English. Get instant answers, make smarter decisions, collaborate more easily—and cut reporting time in half. What questions will you ask? To learn more, check out the show notes or visit www.adverity.com.

00:37:05 Right. Yeah. The behavioral change is tricky, for sure. I've recently launched my own consultancy for bringing cutting-edge things like multi-agent systems, generative AI, exactly as you said, the kind of opportunity that Workhelix, this trillion-dollar opportunity that Workhelix's co-founder Erik brought up where there's so many tasks

in the enterprise today that can be improved or automated fully with a large language model to a high level of accuracy, especially because you can have a second LLM that is double-checking work. You can have cascading LLM systems where for tasks that are more complex or the LLM, an initial cheaper LLM is unsure of, gets passed to a more expensive one. So there's all kinds of tricks that you can have in an organization. It doesn't seem like at this time technical hurdles are preventing enterprise AI adoption. There's plenty of opportunity probably in any organization to be streamlining operations with AI. It is people-

Shaun Johnson: 00:38:14 Yeah, that are tricky. One of the impressive things about the agricultural, the industrial revolution is we went from a world where 95% of our workers of "field workers" were heads down working on these farms and then suddenly you have tractors much more autonomous systems, but the beauty of that time is that you didn't have a mini tractor that everyone had to use. You had to say, "Hey, worker, now you have to use this mini tractor and have all that behavioral change." You just replaced 100 overnight. That's not what AI is, right? AI looks like right now, it's more gradual. We need to go through this phase where there's a co-pilot at first and then there's a human in the loop, and then in some areas of a job, some tasks there can be full automation, but that requires a partnership between the human and the machine, and that's full of friction.

Jon Krohn: 00:39:21 For sure. There's an interesting... I was on a panel at the Open Data Science Conference in Boston last week and got asked a question from the audience about how I thought the world would change in enterprises. If we're talking three years, five years from now, like a relatively short timeframe, how different is life in an enterprise because of AI in all these powerful agent systems that you know, that I know could be adopted in organizations and

they could be so transformative? It's interesting because on that kind of timeline, there are certain kinds of people who think that we'll have artificial general intelligence on that same three to five-year timeline.

00:40:10 I guess my take on that is that in three to five years we'll have superhuman capabilities on more and more narrowly defined tasks. I think that there will still be, because I don't think we'll kind of have the data to be training the world models yet on a three to five-year time scale to have something that is really replacing a human on all kinds of tasks that we do. But regardless, there's going to be really powerful systems in three to five years.

00:40:39 But my answer to the question about what I thought the enterprise would be in three to five years, I was like, "In a lot of ways, it's probably going to be similar. I mean, we're probably still going to be sitting on panels at data science conferences in three to five years complaining about slow paces of change and all the opportunity that there is in enterprises." It's so interesting to see people putting up charts of, "In three to five years, we're going to have AGI. It's going to be transforming absolutely everything." But because of this human friction, I think there's still going to be tons of opportunity on the three to five for your timeframe. What do you think about all that?

Shaun Johnson: 00:41:12 Yeah, I think there's clearly going to be a distribution. Some companies are going to be blazing a path and some companies are not going to have improved that much. I'm always surprised when I'm talking to enterprise leaders, I was talking to one a couple of weeks ago and they were telling me that they run an organization of 5,000 and it's a fairly... It is not a well-known brand. I can call it a very middle of the road brand. Not a lot of people have heard of this brand and so not the most AI-native forward tech leaning, but they said that their team of 5,000 in two to

three years would be reduced to 3,000. I thought that was significant. So I asked, "How could that be?"

00:42:00 As an engineering, there's a few thousand, roughly 30, 40% that are doing patches and lower level work that they feel like AI can handle sufficiently, and so it's just a matter of transitioning the team. So it is going to be a spectrum. Some folks are, like I said, they're going to blaze the path and I think we're all going to be like, "Whoa." Then there's going to be folks where it's like you said, nothing's changed much at all.

Jon Krohn: 00:42:28 Yeah, I think that kind of incrementalism, it is that kind of number that you said it was like 5,000 down to 3,000. That's a huge amount of the company. I guess in my head I would assume that's smaller percentages for the most part. But we recently had, again, at the time of recording in recent days or about in the last week, 6,000 employees were let go from Microsoft and reportedly those are mostly software engineers. It's kind of the same thing you were talking about there, patches. There's lots of kinds of work that can now be fully automated that some kinds of software engineers were needed for up until now, and so it is an interesting... What do you think for our listeners out there, our hands-on practitioners who don't want to be amongst those 6,000 that are being let go at any given time in the coming years, what kinds of things do you think hands-on practitioners can be doing to future-proof themselves?

Shaun Johnson: 00:43:31 Yeah, I think there's a couple of things that we think about. I think in the short term, being an advocate for AI is very important, right? The reality is that machine intelligence is growing quite fast. You have this partner now. It's no longer just this passive device that will take your Excel calculations or Google sheet calculations and perform that for you or run your macro. It's something that can help you think and help you come up with new

solutions to a particular problem. I think everyone needs to be focused. For folks that want to know how to stay far ahead, I think that's number one.

00:44:12 Then I think in the longer term, we talk about this quite a bit in AIX, I think it's going to be really important for knowledge workers should they want to continue in their fields and in the enterprise, I think they're going to have to gravitate towards where there are human-to-human touch points, right? So the digital world is being quite transformed right now. Robotics is still pretty, pretty far away. We think two, three major research breakthroughs to be able to have... When our humanoid robot is going to pass the Turing test, right? So for a long, long time, I think you're going to need that human-to-human interaction. So if you are outbound sales, if you're working with clients, if you are directly connected to just other people, not as much internally, but also externally and at least with the higher levels of the organization, I think that's another way to think about building your career.

Jon Krohn: 00:45:20 Yeah, that's something that I come back to is your ability to influence is something that will keep you safe in this AI era. Still continuing on with this kind of enterprise adoption of AI conversation, but going back to Workhelix, they had a press release saying that, "We're in the first inning." This is a baseball analogy, so there's nine innings in a baseball game for our international listeners who wish that we had more cricket analogies perhaps. So this Workhelix press release says that, "We're in the first inning of a decades-long AI transformation," and I absolutely agree with that. That is what we're talking about already, you and me, Shaun, in recent minutes about how we're going to have this incremental change, patches to software being something that's replaceable and in two years it'll be more advanced software tasks,

machine learning tasks that are increasingly automatable.

00:46:23 There's something interesting here to me with enterprises where it can be really tricky to rigorously measure the impact of an AI solution that's brought in that has streamlined something. I'm experiencing this right now with a client that we're automating an aspect of their business with my consultancy using generative AI tooling. I've been having almost daily meetings with my co-founder to make sure that we can come up with a great metric to ensure that we're able to demonstrate the impact of this.

00:46:58 What do you think are the key things in terms of being able to sell an enterprise AI solution to an enterprise? What are the kinds of key things that we need to get into place, especially if it's difficult to measure AI impact? Being able to demonstrate a return on investment is obviously... That seems to me like an ideal, but sometimes it's even tricky to be able to have rigorous metrics that demonstrate operational efficiency improvements. It can be even harder to get an ROI figure that you believe in. So something that comes to mind for me is actually storytelling is a way to get buy-in influence and get those enterprise AI solutions adopted.

Shaun Johnson: 00:47:45 Yeah. Well, you've called it out that we are Workhelix investors, so that's the disclaimer for the audience, but this is one of the best teams at doing just that. I was talking to James, the CEO of Workhelix the other day, and he's having customers come to him and say, "Hey, can you help us come up with metrics that are things we should look at and things we should expect to move?" So the Workhelix team is not just partnering with enterprises to understand what's interesting to them, but also what should be interesting as a particular metric.

00:48:23 Where I think it gets complicated is when you are increasing productivity, but it's not going to move revenue in any substantial way and it just reduces costs, but maybe it's just less clear because the people that are becoming more productive just become a cost center in a different place. What I mean by that is... Let's just take a dentist office. A small dentist office, there's only two or three people in the front desk. If you make them all more productive, they just go do other things. There's so much entropy in that office and how do you measure that? That just becomes really complex and then suddenly you're just a cost center because you've sold some AI in and it's not good for retention.

Shaun Johnson: Because if you're going to come in and make everyone more productive, riffing is really hard. How are you going to get adoption from a team that knows they're going to be riffed? Instead, you want a team that's so understaffed that they're making errors, that they know they could do better. They have to hire. The hiring process is brutal. Just pull the JDs off the table is, I think, the fastest time to value.

Jon Krohn: For those of our listeners who don't know what a riff is, it isn't grabbing your guitar and jamming with your friends, which is fun. Riffs are not fun. They're a reduction in force, in workforce size. Yeah, that's that term. Another term actually that you've used a few times in this episode that I feel like probably a lot of our listeners would understand just based on context or just saying it is AI native, but I'd love you to define that more specifically. Actually, let's just do that first. I got two follow-up questions related to AI native, but do you want to just define what that means to you?

Shaun Johnson: Yeah, I think there's two contexts in which we hear and use AI-nativeness. One is at the product level, and then another one is at the team level. At the product level,

what I think the general community means when they think about AI native was the product built from scratch with AI front and center? What the converse is, is if there was some legacy product and then you tried to bake AI in. Being able to start from scratch, the idea is that every interaction, every user experience is going to have AI front and center. When you're bolting it on, it feels a little bit more kludgy, if you will. Then on the... Go ahead. Yeah.

Jon Krohn: I was going to say that even from the perspective of this podcast, this podcast is nine years old and we release two episodes every week, every single week of the year. So 104 episodes a year. So it's very difficult for me to get ahead a month and be able to completely overhaul operations or something like that. I would love to be starting a podcast right now so that I could be like, "Okay. Our operations from scratch, we're going to have agents everywhere, gen.ai everywhere." There's so much opportunity. If we were doing it in that AI-native way to give people a concrete example, everything would flow, I imagine, so much easier, whereas once you have operations, there's things that work and the cost, the price of overhauling one piece, when you've got to keep the bus on its wheels and going, and I got to get two episodes out every week, it's hard to be like, "Okay. We're going to completely transform this piece because it could have knock-on effects that I don't anticipate and that hinder my ability to be getting episodes out on time."

Shaun Johnson: Precisely. Then at the team level, I think AI-nativeness is more subjective, but I think the very high level question we ask is, do we think this team will be able to unlock applications and user experiences that most other teams cannot, right? If you can do that, that's an AI-native team. It does indeed depend on domain where you're innovating. But an example of this is Richard's team at you.com. They are able to unlock experiences in you that are representative in AI-native team. Same with our

vendor, Perplexity, very focused on the consumer side of things, able to unlock experiences that indeed are representative in AI-native team.

Jon Krohn: Yeah. Yeah. It's kind of interesting to me how in my mind, in the vector database of my brain, you.com and Perplexity are actually quite close to each other in that space. Yeah, it's interesting to think how... Oh, we don't need to get into that. But in terms of opportunity that you see in across verticals, you sit in a position where you must get pitches in all kinds of verticals, AI-native applications from AI-native teams. For our listeners, for people who are thinking about building something new, if you're able to disclose this, what are the kinds of verticals that you think are ripe still today for disruption by AI companies?

Shaun Johnson: Yeah, I think the reality is if you think about this continuous curve between this world of consensus investing and non-consensus investing, consensus being obvious applications that everyone's building towards and non-consensus being applications have a little less obvious, requires some degree of insight, there's opportunities in both worlds, right? There was opportunities in SaaS five years ago. You just had to have some insight into why the world could change in a formidable way and how you are going to change it. It's still going to be true, right? Sometimes I see investors that are like, "Oh, this space is too busy," call it engineering, AI-native engineering applications. Cursor's there. Windsurf's there. Cognition's there. Too many are there. That's not true, right? These companies have grown bigger and bigger and bigger, and that's when they're ready to be disrupted, just like always.

So there is a founder that is working on a vision of the future that might have a wedge that could grow very quickly. In fact, you hear that Cursor and Windsurf to a

degree are being disrupted by Lovable. So, here's a new player that's disrupting, just a couple year old incumbent. So I wouldn't say that even the red ocean is... Everything was red ocean in SaaS five years ago, right? The VC world, in the startup world, and entrepreneurs continue to be innovative and think about great ideas and great startups that will turn into enduring companies, and so I wouldn't steer anyone away.

If you have a quite capable team and you want to go into red ocean with an interesting idea, do that. We're about to invest in a team today or tomorrow that's doing exactly that. If you have a team that's looking at a vertical that not a lot of folks talk about and it's this... I was talking to someone today that's innovating in CPG and a specific area in CPG transformation when it comes to product development, and it's so specific and their background was so specific to that, it's very compelling.

Jon Krohn: Yeah. Yeah. There's a term that you used in there that I don't think I've come across before. You said red ocean. What does that mean?

Shaun Johnson: Oh, blue ocean, red ocean. Red means sharks, a lot of fighting. Blue ocean just means a lot of opportunity.

Jon Krohn: I see, I see, I see. Yeah, I did have that visualization. I guess it means it's a good term. I kind of had that pop up in my head also. There's an interesting... Yeah, go ahead.

Shaun Johnson: I was just going to say, I think there's a book out there called Blue Ocean Opportunities or something like that.

Jon Krohn: There's an interesting... You pronounce it the word or the acronym, S-A-A-S. I always pronounce that SaaS and I feel like pretty much everyone does, but that obviously SaaS is already a word like sassiness, and so it's interesting that... This might be because I'm in New York

and I'm not in Silicon Valley, and so I'm so out of touch with the way, but you pronounce it SaaS. You mean software as a service every time you say SaaS, right?

Shaun Johnson: Indeed.

Jon Krohn: Yeah. Yeah. So, I like that. I'm probably going to adopt now. So listeners, you heard me say this here first, I'm now going to start pronouncing it SaaS-

Shaun Johnson: SaaS.

Jon Krohn: ... to distinguish against sass, sassiness I guess that's a pretty common pronunciation for you on your side.

Shaun Johnson: Yeah. Yeah. I mean, I hear SaaSter quite a lot, not sasster, SaaSter.

Jon Krohn: Yeah. Yeah. Yeah. Nice. So we've gotten through a ton of the topics that I planned for you. There's a trend that's been happening that I'd love to talk about on air that I don't think I've talked about on air before and get your input on. There's been a lot of instances in recent years where big tech companies like Google, Microsoft, it seems like to avoid regulatory issues instead of acquiring startups, they end up...

In fact, there's... I can't remember which kind of acquisition it was, but just this week, the US Justice Department opened up a new antitrust suit into Google for one of these where instead of the entire company being acquired, they acquire a bunch of the talent, including often the executive team, but it seems like they're trying to avoid these kind of antitrust issues where there was this problem up until a couple of years ago of these big tech players getting increasing scrutiny for making more and more acquisitions, making it difficult for competitors, like you said, these kinds of red

ocean players that are trying to disrupt bigger players instead of acquiring them. So now you see this thing happening more and more where companies' leadership joins Microsoft, joins Google, but the whole company isn't. From your perspective, is that changing strategies that startups have? Is that something that you like as a VC, or is that something that actually has the potential to undermine investments that you make?

Shaun Johnson: Yeah. So there are very few, let's call it, maybe single-digit thousands right now of AI practitioners that are at the frontier, right? There aren't many of them and that means that there's a talent disparity. There's massive scarcity, which creates a disparity between different companies, the haves and have nots. And so when that happens, what do you expect to happen? The companies are going to increase salaries, compensation. I heard the other day about a new Google building that came up that's very, very fancy, that's just focused on AI. It's color coded and it's just like it's next level. I think you're seeing compensation skyrocket. I was talking to a frontier engineer the other day who said that they just joined one of the hyperscalers and money's no longer an issue. They don't really think about the price of things anymore, right? It's so competitive.

When you're operating in that kind of environment, you'll do anything you can to buy up talent to differentiate your business, because what you do now, it all compounds. What you're able to accomplish this week and then next week and the following week is compounding. If you lose your edge in three to five years, you're going to see it in the market. So, what does that mean? I'll tell you, I think there are pros and cons as a venture firm. I think the pro is you can talk to have a conversation with founders, founders that have raised capital, that have enough capital to get through 24 months and they're operating at a healthy pace. But you can say you need to recognize

that you have a very large opportunity cost here, right? You could be at a big company, which they don't want to be at because they're starting a company.

They want to start their thing, but you could be at a big company making a ton of cash right now, and instead you've decided to go the startup route. You need to go, go, go, right? Make sure you're blazing here because the opportunity cost is so large. Now, the counter to that is sometimes they might just be like, "Hey, I'm going to go and get bought and go work at those large companies." So in that case, you may lose a founder or a team to one of these companies, but on the other hand, it provides downside protection, right? If you're an investor and you're investing in AI-native teams, you can be sure that you're going to get your money back or some large fraction on the dollar through the acquisition. We had one team that went from fully operating to acquired by Google in three days. It was so fast and we were able to get all of our... Maybe even make a little money on that.

Jon Krohn: That's a great perspective. It is pretty interesting to think about these kinds of... You talked about the single-digit thousand AI experts at the frontier. I don't know if you'd know the answer to this question, but I'd love to hear your two cents if you do have some thoughts here. For our listeners out there who are thinking, "I'd love to be amongst those few thousand people," what are the kinds of things that you can do to get there? What do you need to learn or what organization do you need to be a part of to get there?

Shaun Johnson: Yeah, I think there's different paths for people depending on where they currently are in their career, right? If your background is an engineer that's building and you have an inclination for math and some research, you might think about, how do you work your way into fair at Facebook or into DeepMind? Then if you're just starting

off your academic career or studying CS at the undergraduate level, the master's level, you might consider a PhD in AI. That is a fairly strong path, especially from universities like, call it, CMU or Berkeley or Stanford or NYU or MIT Toronto. There are bleeding edge AI. I heard the other day that Rice actually university in Texas has 100 GPU cluster for their students. So even some of the non-obvious schools are taking AI quite seriously. So that's a path for a student.

Jon Krohn: Actually, I saw on social media just yesterday at the time of recording a funny post. I won't be able to find it in time for it to be of interest to listeners, but I'll try to maybe find it for the show notes and put it in there. Rice University, they have a new undergraduate AI course and their branding for it was one of the only programs preparing you for AI in the real world, in the world. The person who shared it on LinkedIn was like, "This is shameful to say something like that, that there's one of a few degree programs in AI." They are popping up all over the place. But yeah, the kinds of institutions that you mentioned, they're doing PhDs in AI to obviously be at the cutting edge. You're not talking about an undergrad in AI. It's a PhD in AI from one of these top organizations.

But yeah, I also love how you have the kind of alternative route there. Obviously you can learn anything at any age. So if you already are highly technical, you like doing research, like you said, you're highly numerate, you like doing math, then you can see what's in the job descriptions for a DeepMind AI engineer and start pursuing that.

Shaun Johnson: If you found that your bedtime reading turns out to be linear algebra, you should try it out.

Jon Krohn: Exactly. I've never been in a situation where I am hiring people of that kind of expertise. I've certainly never been,

I'd love to be, in a position that I was hiring from the single-digit thousand kind of AI experts in the world, but I did once have an interview. I interviewed someone where I was told by the person that introduced me that they do read linear algebra in bed, and I was really excited to meet them, but it was a really disappointing interview. They weren't that strong, so I don't know.

Anyway, I've got one last question for you before I get into my final... As I mentioned to you before we started recording that I always ask for a book recommendation and I also always ask our guests how people should be following them. But my last deep question before we get there, so far in this episode, we've been leveraging your knowledge, but I haven't talked much about your particular background, which is the way that we do things on this show. Most podcasts, they have people just talk about their career and how they got to where they are. But I want to get to what are you doing now and what's really exciting, which is obviously where we're focused on the episode. But sometimes at the end, I want to go back a little and I'm going to do that now with one last question, one kind of deep question.

So you have a multidisciplinary background spanning engineering business and sustainability. How do you leverage that background to help founders that you back integrate customer empathy and long-term impact considerations into AI products? It seems like that's something that's important to you because otherwise founders might just end up prioritizing technical performance above all else.

Shaun Johnson: Yeah/ yeah. I don't know. Well, maybe this two birds with one stone, we'll see. But there's this book called Range, right? I mean, I think it broadly argues that if you see problems in a lot of different spaces, then you can bring a broader skillset to founders when you're looking at their,

whether it be product market fit or hiring or any numerous things they're looking at and just be more resourceful. So especially coming through my career from CS to EE to business, to even within CS and EE, from being an individual contributor to being a manager, to working more on the marketing side to all these different things, I think I've just seen a lot, right?

Now, the beauty of the AIX is I'm seeing way more even every day because I get to talk to so many different founders and pick their brains and get to know them pretty well and support our portfolio and see how they're doing. So I like to think that every day I get better and I improve my own skillset and can be even more valuable to founders. So that's a little bit about how I think about it broadly.

Jon Krohn: Nice. Yeah, you got the book recommendation in there as well. Nice. Was that the book that you wanted to recommend or do you want kind of-

Shaun Johnson: Yeah, I was thinking about a book that I've read a long time ago and continue to read that I quite like. The Art of Happiness by Dalai Lama I think is quite a good book to just level set and continue to be appreciative every day of the time that we're living in and the people and the innovation we're surrounded by is staggering. What an amazing time. Then books like Range. There's a new book called AI Valley that came out talking about Silicon Valley and this time and those are all good readings.

Jon Krohn: Nice, great recommendations there. I haven't read that particular book by the Dalai Lama, but I read his... I think it's called an Open Heart some years ago. That was a good one. And I got to say, you do seem like an unusually happy person, so it's probably just that one book.

Shaun Johnson: I have a four-year-old that is getting to the point where he knows better than I, so I'm always very humble.

Jon Krohn: Nice. All right. Shaun, this has been an amazing episode. I've learned so much from you. I'm sure a lot of our audience has as well. What's the best way for them to follow you after this episode and get more of your thoughts?

Shaun Johnson: Yeah, I'm on LinkedIn. You can find me pretty easy. I'm at Twitter, X, Shaun B. Johnson. My email is just s@aixventures.com. S, it's just first letter, my name, S as in Shaun, @aixventures.com.

Jon Krohn: Nice. Thank you so much again, Shaun, for taking the time. Yeah, maybe we'll catch up again with you in a few years and see if we can complain again about how enterprise AI adoption is going too slowly at that time.

Shaun Johnson: Thanks so much, Jon.

Jon Krohn: What an honor to have Shaun Johnson in today's episode. In it, the iconic AI investor covered how AIX Ventures unique model combines AI practitioners like Richard Socher and Chris Manning with traditional VCs to evaluate both technical depth and market-savvy in founding teams. He said the most effective enterprise AI adoption strategy is to target open job requisitions. Instead of augmenting current workers, AI can fill positions companies are already trying to hire for. He opined that only single-digit thousands of AI experts operate at the true frontier, creating massive talent scarcity and opportunity costs that drive both high compensation and talent acquisitions.

He talked about how both consensus applications in competitive spaces and non-consensus bets in overlooked markets can both succeed execution speed and the ability

to pivot matter more than finding a perfectly empty market. He also talked about how this trend of major tech companies are increasingly doing talent acquisitions rather than full company acquisitions to avoid regulatory scrutiny while also still securing scarce AI expertise. Finally, Shaun talked about how knowledge workers should focus on human-to-human touchpoints and building influence to remain valuable as AI transforms digital work.

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 Shaun's social media profiles, as well as my own at superdatascience.com/895. Thanks to everybody on the SuperDataScience podcast team, our podcast manager, Sonja Brajovic, media editor, Mario Pombo, Nathan Daly and Natalie Ziajski on partnerships, our researcher, Serg Masis, our writer, Dr. Zara Karschay, and yes, our founder, Kirill Eremenko. Thanks to all of them for producing another excellent episode for us today. For enabling that super team to create this free podcast for you, we are deeply grateful to our sponsors. You, listener you, yes, can support this show by checking out our sponsor's links, which are in the show notes.

If you yourself are interested in sponsoring an episode, you can find out how to do that at jonkrohn.com/podcast. Otherwise, you can help us out by sharing this episode with folks who would love to listen to it as well, reviewing the show on whatever platform you listen to it on, subscribe. Feel free to edit videos into shorts, but most importantly, just keep tuning in. I'm so grateful to have you listening and I hope I can continue to make episodes you love for years and years to come. Until next time, keep on rocking it out there and I'm looking forward to enjoying another round of the SuperDataScience Podcast with you very soon.